arXiv:1404.0667v3 [math.DS] 12 Oct 2015

ATLAS: A geometric approach to learning high-dimensional stochastic systems near manifolds

Miles Crosskey[1], Mauro Maggioni[1,2,3]

Department of Mathematics[(1)], Electrical and Computer Engineering[(2)] and Computer Science[(3)], Duke University, Durham, NC, 27708

October 13, 2015

Abstract

When simulating multiscale stochastic differential equations (SDEs) in high-dimensions, separation of timescales, stochastic noise and high-dimensionality can make simulations prohibitively expensive. The computational cost is dictated by microscale properties and interactions of many variables, while the behavior of interest often occurs at the macroscale level and at large time scales, often characterized by few important, but unknown, degrees of freedom. For many problems bridging the gap between the microscale and macroscale by direct simulation is computationally infeasible. In this work we propose a novel approach to automatically learn a reduced model with an associated fast macroscale simulator. Our unsupervised learning algorithm uses short parallelizable microscale simulations to learn provably accurate macroscale SDE models, which are continuous in space and time. The learning algorithm takes as input: the microscale simulator, a local distance function, and a homogenization spatial or temporal scale, which is the smallest time scale of interest in the reduced system. The learned macroscale model can then be used for fast computation and storage of long simulations. We prove guarantees that related the number of short paths requested from the microscale simulator to the accuracy of the learned macroscale simulator. We discuss various examples, both low- and high-dimensional, as well as results about the accuracy of the fast simulators we construct, and its dependency on the number of short paths requested from the microscale simulator.

1 Introduction

High-dimensional dynamical systems arise in a wide variety of applications, from the study of macromolecules in biology to finance, to multi-agent systems, to climate modeling. In many cases these systems are stochastic by nature, or are well-approximated by stochastic processes, for example as a consequence of slow-fast scale phenomena in the system. Simulations typically require significant amounts of computation, for several reasons. First of all each time step of the numerical scheme is often expensive because of the large dimensionality of the space, and the large number of interactions that need to be computed. Secondly, fast timescales and/or stochasticity may force each time step to be extremely small in order to have the requested accuracy. Finally, large-time behavior of the system may be dominated by rare transition events between stable regions, requiring very long paths to understand large-time dynamics. A large amount of research spanning multiple fields tackles the problems above.

Suppose we are given a high dimensional stochastic simulator, and we are interested in the large-time behavior of the system, but are faced with the problem of prohibitively expensive costs to run long simulations. What could be computable in a highly parallel fashion is an ensemble of short paths [49]. We therefore ask: what can be learned from ensembles of short paths? Several crucial problems to be addressed include: where in state space such short paths should be started? how many paths should be run locally? for how long? how does the local accuracy depend on these parameters? and once these local paths are constructed,

and perhaps local simulators approximating the local dynamics are constructed, how can they be stitched together to produce a global simulation scheme? What can be guaranteed about the accuracy of such a global scheme, for large times? Some examples, among many, in this direction are Markov State Models [39, 8, 49], milestoning [18], and several other techniques (e.g. [37, 27, 15, 47, 29, 53, 21, 29, 3, 28] and references therein). Many of these methods are based on discretizations of the state space into regions, and measure transitions between such regions, others on biasing the potential to speed up exploration, yet others estimate local statistics and use them to coarse grain the system. Our method is related to some of these existing approaches, but uniquely combines them with ideas from machine learning, stochastic processes and dimension reduction, and introduces several novel key elements that combined lead to an accurate yet low-dimensional estimate of the generator of the diffusion process (rather than estimating discretized transition densities), with guarantees on the large time accuracy of the simulator we construct. Our construction can be seen as a higher order approach compared to Markov state models, since we fit a linear reduced model, continuous in space and time, to each region, and smoothly glue these models together. Because of this, we are also able to approximate the original dynamics by a process which is continuous in time and space.

The philosophy of reducing a high-dimensional system to a low-dimensional surrogate is well-established as enabling the simulation of complex, large, high-dimensional systems, and more methods have been proposed than we can possibly discuss here. These include model reduction [35, 3, 28], homogenization of PDE's [25, 22], coarse-grained dynamics of high-dimensional systems [21, 29], and multiscale modeling [1, 47, 29, 53]. We refer the reader to [57] for a summary of the motivations and applications of several of these techniques, and to the references therein.

We take concepts from *manifold learning* [19, 9, 59, 46] in order to learn an underlying low-dimensional manifold around which most trajectories concentrate with high probability. We approximate the macroscale manifold with linear low-dimensional subspaces locally, which we call charts. These charts enable us to learn local properties of the system in low dimensional Euclidean space. Geometric Multi-Resolution Analysis (GMRA), introduced in [19] uses this concept to approximate high dimensional distributions on manifolds. These techniques perform the *model reduction* step, mapping the high-dimensional system from $\mathbb{R}^D$ down to d-dimensions, yielding a small set of coordinates describing the effective small number of degrees of freedom of the system.

We combine this dimension reduction step with *homogenization theory* [20, 22, 53] and learn a local low-dimensional approximation of the system at a certain time-(or space-)scale t_0. We note that this approximation is not necessarily accurate at timescales shorter than t_0. Short-time events may be complex, high-dimensional, highly stochastic or deterministic but chaotic, and we do not seek a simulator reproducing these fine-scale behaviors. We are interested though in the "net effect" and implications of these phenomena at timescale t_0 and larger. Locally we fit a simple reduced system, e.g. a constant coefficient SDE to each chart. If the macroscale simulator is well approximated by a smooth SDE, then constant coefficient SDEs will approximate the system well locally. This smooth SDE is approximating the original simulator above a certain timescale.

In order to obtain a global simulator, we add a last crucial ingredient: we construct an approximation to the transition maps between charts, generating a numerical approximation to a manifold *atlas*. Learning such transition maps between charts is necessary to allow us to smoothly combine simulators on distinct charts into one global simulator on the atlas. The simulator we construct we then call the ATLAS, and we show that under appropriate conditions it captures long term statistics of the dynamics of the original system.

Finally, we note that accurate samples from the stationary distribution is a valuable tool in studying many dynamical systems. Reduced large-time models for complex high dimensional dynamical systems is sometimes obtained using so called "reaction coordinates", a set of global low dimensional coordinates describing the important states of the system (see e.g. [36, 43, 10, 24, 13, 52, 14, 42, 50, 4, 18] and references therein). Several of the techniques (but not all, e.g. notably the string method of [54, 16, 44]) need samples from long simulations of the system and/or the stationary distribution. For example diffusion maps and their extensions to the study of molecular dynamics data (see [12, 11, 45]) require many samples from the stationary distribution to guarantee accuracy. These reaction coordinates allow further analysis of dynamical systems

by easily identifying stable states, and, most importantly, transitions paths between such states and reaction coordinates parametrizing such transitions. In fact, part of the motivation for this work was observing that the slowest part of running diffusion maps on such complicated high dimensional systems was obtaining the samples from the stationary distribution.

The paper is organized as follows: in section 2 we describe at high level our construction, algorithm, and informally state the main result on the accuracy of the ATLAS for large times; then we illustrate the algorithm on simple examples. In section 3 we discuss the algorithm in detail. In section 4 we state and prove our main result. In section 5 we present a wide range of examples. We conclude with a discussion in section 6.

2 Construction and Main Results

The geometric assumption underlying our construction is that the dynamics of the Markovian stochastic dynamical system of interest $(Y_t)_{t\geq 0}$ in $\mathbb{R}^D$ is concentrated on or near an intrinsically low-dimensional manifold $\mathcal{M}$ of dimension d, with $d \ll D$. We refer to $\mathcal{M}$ as the *effective state space* of the system, as opposed to the full state space $\mathbb{R}^D$. This type of model may be appropriate in a wide variety of situations:

(i) the system has d degrees of freedom, and is therefore constrained (under suitable smoothness assumptions) to a d-dimensional manifold $\mathcal{M}$;

(ii) as in (i), but possibly with small deterministic or stochastic violations of those constraints (perhaps at a fast scale), but such that the trajectories stay close to $\mathcal{M}$ at all times.

In these cases it makes sense to approximate $\mathcal{M}$ by an efficient low-dimensional approximation $\mathcal{A}$, such as a union of d-dimensional linear affine sets (charts) [2, 33], and the dynamics of Y_t by surrogate dynamics on the atlas $\mathcal{A}$. Learning dynamics on $\mathcal{A}$ reduces the problem from learning a high-dimensional global simulator to a low-dimensional local simulator, together with appropriate transitions between local simulators in different charts. We also gain computational efficiency by using the structure $\mathcal{A}$: long paths may be more quickly stored and simulated in lower dimensions. We will make assumptions about the geometry of $\mathcal{M}$ and on the underlying macroscale simulator, in order to prove large time accuracy results for the ATLAS. We expect this approach to be valid much more generally, and this hope is supported by our numerical experiments.

While in this paper we consider a special class of stochastic dynamical systems, those well-approximated by low-dimensional SDEs such as those leading to advection-diffusion equations along a manifold, the framework can be significantly extended, as we briefly discuss later in section 6, and this will be subject of future work.

2.1 Main Ideas and Steps

Our construction takes as input:

- a dense enough sample of $\mathcal{M}$, or a way of sampling $\mathcal{M}$ in a rather uniform way (both of these statements will be quantified later, see section 2.1.1);
- a simulator $\mathcal{S}$ for the stochastic dynamical system $(Y_t)_{t\geq 0}$, which may be started upon request at any specified initial condition and run for a specified amount of time;
- a distance function ρ to be used to for measuring distances between pairs of data points returned by the simulator;
- a spatial homogenization parameter δ;
- the dimension d of the effective state space, and a confidence parameter τ.

We note here that the homogenization scale δ can also be given as a temporal scale t_0, and the two are related by natural scalings in the underlying dynamical system. Given a time t_0, running paths of length t_0 and examining the average distance traveled by such paths reveals a corresponding natural spatial scale $\delta(t_0)$ (in fact this is done in example 5.5). Inversely, given δ, one could choose t_0 so that the average distance traveled by paths is approximiately δ. We later discuss the accuracy of the simulator, which is a function of the parameter δ.

We remark that while d is here considered as a parameter for the algorithm, in fact there is a lot of work on estimating the intrinsic dimension of high-dimensional data sets that would be applicable here. In particular, the Multiscale SVD techniques of [31, 32] have strong guarantees, are robust with respect to noise, and are computationally efficient. See also [33] for finite sample guarantees on the approximation of manifolds by local affine approximate tangent spaces. We will mention again the problem of estimating d when we construct the local charts in section 3.2.

The confidence parameter τ sets the probability of success of the algorithm (at least $1 - 2e^{-\tau^2}$), and is related to the number of sample paths one must use to approximate the local parameters of the simulator.

Our construction then proceeds in a few steps:

(i) **net construction**: find a well-distributed set of points $\Gamma = \{y_k\}$ in $\mathcal{M}$, having a granularity parameter δ, the finest resolution of interest;

(ii) **learning the atlas**: learn local charts C_k near y_k obtained by mapping $\mathcal{M}$ locally to d-dimensional Euclidean domains, and learn transition maps for changing coordinates between a nearby pair of charts;

(iii) **learning the simulator**: run $p = p(\delta, \tau)$ paths for time $t_0 = t_0(\delta)$ from each y_k and map them to the coordinate chart C_k. Use these low dimensional representations to estimate a simple simulator on each chart C_k.

2.1.1 Net construction

The first stage is to produce a δ-net $\Gamma = \{y_k\}$, which is a set of points $\{y_k\}$ in $\mathcal{M}$ such that no two points are closer than δ, and every point in $\mathcal{M}$ is at least δ close to some y_k. With abuse of notation, the range of k will also be denoted by Γ, so we may also write the net as $\{y_k\}_{k\in\Gamma}$. We say that two points y_k and y_j are connected, or $k \sim j$, if y_k and y_j are within 2δ. We shall construct a reduced simulator at each node and the neighboring connections will determine the switching between simulators at adjacent nodes. See section 3.1 for the details.

In real world examples, the space $\mathcal{M}$ may be unknown. In this case we assume that we have the ability to sample from $\mathcal{M}$, and we generate enough samples $\{x_i\} \subset \mathcal{M}$ such that balls of radius $r \ll \delta$ cover $\mathcal{M}$. This first round of sampling should ideally have the following properties: it can be generated by a fast exploration method (e.g. see the recent work [61] for molecular dynamics, and references therein - this problem by itself is subject of much research); its samples do not require a significant number of calls to the simulator, or long runs of the simulators; different points may be sampled independently so that the process may be parallelized. We can then downsample these $\{x_i\}$ to obtain the desired net Γ. These considerations depend on the sampling measure. If this is simply the canonical volume measure on $\mathcal{M}$, it is easy to see that $O(\delta^{-d}\log(1/\delta))$ samples suffice, with high-probability, to obtain the desired δ-net. Similarly for measures that have uniformly lower and upper bounded density with respect to the volume measure. We remark that this sampling is therefore independent of the dynamics of interest, has in general nothing to do with the stationary measure of the process we will seek to approximate: it may therefore be significantly easier to construct this sampling mechanism rather than one adapted to the dynamics.

Finally, it is important to remark that the algorithm we present is easily modified to run in exploratory mode: the fast simulator runs on the currently explored region of space, and whenever configurations outside the explored region of space are encountered (an event that is quickly detectable using the data structures we employ), new charts and local simulators may be added on-the-fly with minimal computational work. This is subject of current work and will be detailed in a forthcoming publication.

2.1.2 Learning charts and corresponding maps

The first step in learning the charts is to generate a set of landmarks $A_k \subset \mathcal{M}$ for each y_k in the net $\Gamma \subset \mathcal{M}$: a set of points well-spread in d directions on $\mathcal{M}$, at distance about δ from y_k. In our setting where $\mathcal{M}$ is d-dimensional, we can sample $p \geq d$ paths from the simulator of Y_t, starting at y_k and run until time $t_0 = t_0(\delta)$. As long as the diffusion Y_t is nondegenerate on the tangent plane, the projections of the end points of these paths will span the tangent space to $\mathcal{M}$ at y_k.

Next we learn a mapping Φ_k from a neighborhood of y_k to a coordinate chart $C_k \subseteq \mathbb{R}^d$ for each y_k. In order that neighboring coordinate charts overlap on a region of size δ, we learn Φ_k from $L_k = \bigcup_{j \sim k} A_j$, the union of neighboring landmarks. In this way Φ_k will be defined as a map from $B_{2\delta}(y_k) \subseteq \mathcal{M}$ to $C_k \subseteq \mathbb{R}^d$. The overlap between neighboring charts will allow us to smoothly transition the simulator from one chart to the next. Each mapping Φ_k is constructed using Landmark Multi-Dimensional Scaling (LMDS) on L_k, minimizing distortion of pairwise distances between the landmarks L_k (see section 3.2).

For any $k \sim j$, L_k and L_j have the landmarks $A_k \cup A_j$ in common; thus the charts C_k and C_j overlap on $A_k \cup A_j$. These landmarks span the local charts, and are the points used to learn the transition maps between neighboring charts. The affine transition map $S_{k,j}$ is chosen as the "best" linear mapping from $\Phi_k(A_k \cup A_j)$ to $\Phi_j(A_k \cup A_j)$ described in section 3.3. Figure 1 shows a cartoon version of the points used to learn the atlas.

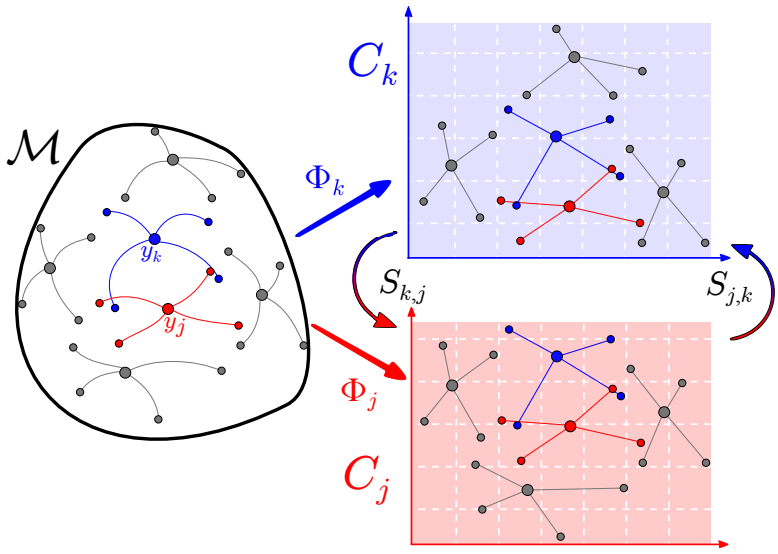

Figure 1: This figure depicts $m = 4$ samples per net point being used to learn the charts. Large circles represent net points (or projections of net points) and small circles represent path end points (or projections of the path endpoints). The LMDS mappings Φ_k, Φ_j use all the circles to learn the chart, while the transition maps $S_{k,j}, S_{j,k}$ use only the colored circles.

2.1.3 Learning the Simulator

Once the charts are known, we learn an approximation to the simulator on each chart. For each $y_k \in \Gamma$, we run $p = p(\delta, \tau)$ paths via the original simulator $\mathcal{S}$ up to time $t_0 = t_0(\delta)$ starting from y_k. Next we project the samples to C_k in order to estimate local simulation parameters. In this paper we use constant coefficient SDEs to model the simulator on each chart:

$$d\bar{X}_t = \bar{b}_k dt + \bar{\sigma}_k dB_t \,, \tag{1}$$

for some $\bar{b}_k \in \mathbb{R}^d$ and some positive definite $\bar{\sigma}_k \in \mathbb{R}^{d\times d}$. The solution to this constant coefficient SDE is a Gaussian with mean $\bar{b}_k\, t_0$ and covariance $\bar{\sigma}_k \bar{\sigma}_k^T\, t_0$. Therefore, we estimate $\bar{b}_k$ and $\bar{\sigma}_k$ by imposing that these statistics match the sample mean and sample covariance of the endpoints of the p paths run with $\mathcal{S}$. Finite sample bounds for these empirical values determine how large p should be in order to achieve a desired accuracy (δ) with the requested confidence (τ).

This step is trivially parallelizable, both in k (the chart in which the learning takes place) and within each chart (each of the p paths may be run independently). At the end of this process we have obtained the family of parameters $(\bar{b}_k, \bar{\sigma}_k)_{k\in\Gamma}$ for a family of simulators $(\widehat{\mathcal{S}}_k)_{k\in\Gamma}$.

The local simulators $(\widehat{\mathcal{S}}_k)_{k\in\Gamma}$ are extended to a global simulator $\widehat{\mathcal{S}}$ on $\mathcal{A}$ using the transition maps between charts. This is done by alternating between steps from $(\widehat{\mathcal{S}}_k)_{k\in\Gamma}$, and transition map operations – this is somewhat delicate, and detailed in section 3.5.

The choice of the local SDE's and the estimator of its parameters is one of the simplest possible, however we will see a collection of these simple simulators combine to reproduce much more complicated systems. Naturally the ideas may be extended to richer families of local SDE's, for which appropriate estimators based on the statistics of local trajectories may be constructed: this is subject of current research.

2.2 Theoretical guarantees

We present here a simplified version of the main result, Theorem 4.1. Suppose the given stochastic dynamical system Y_t is driven by an SDE on a d-dimensional manifold $\mathcal{M}$ with volume measure μ of the form

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t \tag{2}$$

with b, σ Lipschitz functions, and σ uniformly nondegenerate on the tangent bundle $T(\mathcal{M})$. Let q be the density of the stationary distribution of Y_t on $\mathcal{M}$, and $\hat{q}$ be the density of a probability measure on $\mathcal{A}$ defined later in equation (16) and computed by running the ATLAS for large time. Let G be the inverse mapping from $\mathcal{A}$ to $\mathcal{M}$ defined in section 4.3. If the number of sample paths collected at each of $O(\delta^{-d})$ starting points is at least $O((d+\tau^2)/\delta^4)$, then with probability at least $1 - 2e^{-\tau^2}$

$$||q - G_*\widehat{q}||_{L^1(\mathcal{M})} < c\delta \ln(1/\delta) \tag{3}$$

for some constant c depending on geometric properties of $\mathcal{M}$, the Lipschitz constants of the drift b and diffusion σ, and the lower bound on singular values of σ along the tangent plane. Here $G_*\widehat{q}$ is the pushforward of the measure $\widehat{q}$ from $\mathcal{A}$ to $\mathcal{M}$ as defined in equation 17.

One can think of Y_t as the underlying homogenized system which we are trying to learn. This result guarantees that the ATLAS process learned only from short paths of Y_t, actually behaves closely to Y_t for large times. Note that if the microscale simulator does not satisfy these conditions, it is possible the system is well-approximated by a macroscale simulator of the form (2) satisfying the conditions of the theorem on the timescale t_0 (the time sample paths are run) and above; in this case the error in approximating the original simulator by Y_t is simply added to the right hand side of (3).

Our results could be re-interpreted in the context of adaptive MCMC as follows. Assume we wish to sample from a probability distribution q on a d-dimensional manifold $\mathcal{M}$, that is the stationary distribution of a process Y_t as above (with the assumptions stated in Theorem 4.1), then if we have access to a local simulator of Y_t, we can construct an efficient sampler for an approximation to q. Contrary to Riemann manifold Hamiltonian Monte Carlo [23] we do not need to know the parameter space of the underlying statistical model, which would correspond to a parametrization of $\mathcal{M}$, but we learn it through many short simulations of the dynamics, nor do we need sophisticated numerical integrators. These ideas are being developed further in a forthcoming publication.

2.3 Examples

Here we present some examples showcasing the usefulness of the ATLAS. The examples shown here have Brownian motion in a potential well, although the theorem guarantees accuracy for any simulator of the

form (2). Further examples will be discussed in section 5.

2.3.1 Brownian Motion on a Manifold

Given a d-dimensional smooth compact manifold $\mathcal{M}$, one may construct the potential

$$U_\varepsilon(x) = \frac{1}{\varepsilon}\text{dist}(x, \mathcal{M})^2$$

and consider the Itô diffusion in $\mathbb{R}^D$ given by

$$dY_t = -\nabla U_\varepsilon dt + dB_t \tag{4}$$

If one simulates (4) numerically for small ε, the timesteps must be at least as small as $O(\varepsilon)$ (using Euler-Maruyama, since we only need the weak convergence of the scheme). We can view this thin potential around the manifold as our microscale interactions which forces our choice of timestep. What we are interested in is the macroscale behavior determined by the manifold $\mathcal{M}$.

For $\varepsilon \to 0$ this process converges to the canonical Brownian motion on the manifold $\mathcal{M}$ [17]. For ε sufficiently small (compared to the curvature of $\mathcal{M}$) Y_t is well-approximated locally by (the low-dimensional) Brownian motion on $\mathcal{M}$, and the stationary distribution of Y_t is close to that of Brownian motion on $\mathcal{M}$. Our results apply to this setting, yielding an efficient d-dimensional simulator for Y_t, without a priori knowledge of $\mathcal{M}$.

2.3.2 One Dimensional Example

In this numerical example, we start with Brownian motion in a simple double well, and add a high frequency term to the potential to get $U(x)$:

$$U(x) = 16x^2(x-1)^2 + \frac{1}{6}\cos(A\pi x) \tag{5}$$

with $A = 100$. The high frequency term gives the Lipschitz constant $L \sim 10^2$, forcing the forward Euler scheme to use time steps on the order of $L^{-2} \sim 10^{-4}$ in order to just achieve stability (using a higher order method would not solve these problems as higher derivatives of U will be even larger; implicit schemes would allow for a larger step size, at the expense of larger computational complexity for each step). The first term in U is much smoother, and homogenization theory (e.g. [40] and references therein) suggests that the system is well-approximated by a smoother system with Lipschitz constant $l \sim 10$ or less (determined by the quartic term in U), at least for A going to ∞ (and suitable renormalization). Our construction yields an approximation of the given system, for A fixed, at a spatial scale larger than δ (corresponding to a timescale larger than δ^2). This is the target system we wish to approximate. Running ATLAS with $\delta = 0.1 \sim l^{-1}$, we obtain a smoothed version of the potential homogenizing the high frequency term (see Figure 2).

The ATLAS takes time steps which are over 10^2 times larger than the original system, and thus long paths can be simulated 10^2 times faster, as if the system had a potential smooth at scale $O(1)$. Note that increasing the frequency of the oscillating term (thereby increasing L) does not affect the speed of the ATLAS, but only the speed of constructing the ATLAS (since $\mathcal{S}$ would have to use smaller time steps). This means that our algorithm allows for a decoupling of the microscale complexity (which is handled in parallel) from the macroscale complexity. A histogram of the approximate stationary distributions are shown in Figure 3 comparing the ATLAS and the original system. See section 5.2.2 for more details about the experiment, and Figures 13, 14 for the errors in true effective potential vs. estimated effective potential, and the error in approximating the time evolution of the original system by the ATLAS for a multiscale choice of times, as well as the different in transition rates between the wells.

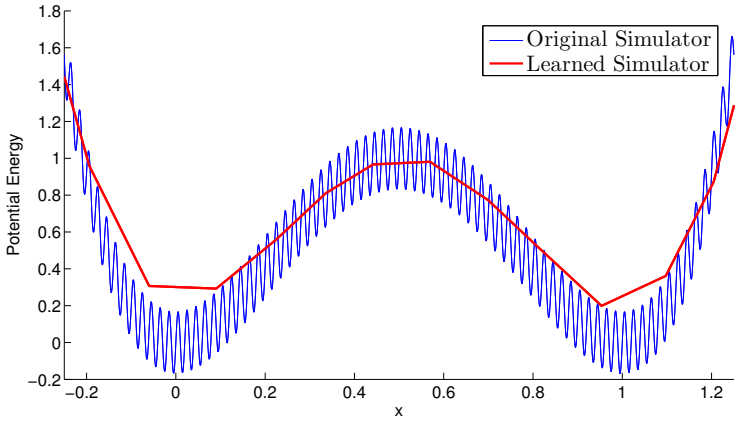

Figure 2: Original stiff potential U (shown in blue), and effective potential $\widehat{U}$ (shown in red) for the ATLAS, learned from short trajectories of the original simulator.

3 ATLAS Algorithm

In this section we present the algorithm in detail, since the main result will state properties of the output of the algorithm itself. Pseudo-code is presented in figure 4. We start by discussing the algorithms used during the learning phase. We discuss afterwards the full details of the simulator learning phase and the simulation phase. The algorithm has access to a simulator $\mathcal{S}$ of the original system, and takes as input several parameters:

- δ: We will assume this parameter is given to us and represents the homogenization scale, and it will be related to the desired accuracy of the simulator via (19) in Theorem 4.1.
- $\{x_j\}$: A set of points on $\mathcal{M}$ that is dense enough that a δ-net for $\mathcal{M}$ may be extracted from it. Alternatively, a way of sampling points on $\mathcal{M}$ with respect to a measure comparable to volume measure on $\mathcal{M}$.
- t_0: This represents the time short paths will be simulated for. In most examples in this paper, we choose $t_0 = \delta^2$. In practice, one should choose t_0 so that sample paths at time t_0 are an expected distance δ from the starting location.
- m: The number of landmarks for each net point for learning the chart and transition maps. m should be at least d, and we choose it of order d to minimize sampling and computational complexity.
- p: The number of sample paths computed for each point in the net. p should be $O(\delta^{-4})$.
- Δt: time step of the ATLAS. It should be $O(\delta/\ln(1/\delta))$. In the examples we used $\delta/5$.

These choices of parameters are informed by the results and proofs in section 4. We will see that for these choices of parameters, the ATLAS produces paths whose stationary distribution has error $O(\delta \ln(1/\delta))$.

3.1 Net construction

In a metric space $(\mathcal{M}, \rho)$ we define a δ-net of points as follows:

Definition 3.1 (δ-net)**.** *A δ-net for a metric space $(\mathcal{M}, \rho)$ is a set of points $\{y_k\}_{k\in\Gamma}$ such that $\forall k_1, k_2 \in \Gamma$ $\rho(y_{k_1}, y_{k_2}) \geq \delta$, and at the same time $\forall x \in \mathcal{M}\, \exists k \in \Gamma\, \rho(x, y_k) \leq \delta$.*

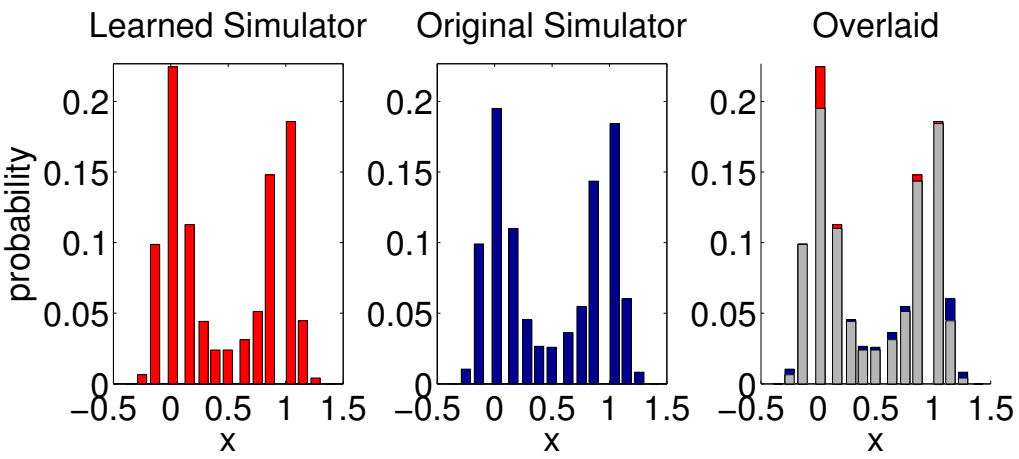

Figure 3: Comparison between the (approximate) stationary distributions between the original simulator and the ATLAS with 10^5 samples for example 2.3.2 (for this number of samples the error is expected to be $O(\sqrt{10^{-5}}) \ll \delta$).

In view of our purposes, the first property ensures that the net points are not too close together: this is essential so we do not waste time exploring regions of the space that we have already explored, do not construct many more local simulators than needed, and do not switch between charts significantly more often than necessary. The second property ensures that $\{B_\delta(y_k)\}_{k\in\Gamma}$ is a cover for $\mathcal{M}$, guaranteeing that we explore the whole space $\mathcal{M}$. We will connect nearby net points: if $d(y_{k_1}, y_{k_2}) \leq 2\delta$ we say that y_{k_1} and y_{k_2} are neighbors and we write $k_1 \sim k_2$.

3.1.1 Computational cost

Algorithms for efficiently constructing δ-nets in metric spaces satisfying a doubling condition exist and are non-trivial, for example by constructing a data structure called cover trees (see [5]), which run in $O(C^d n \log(n) D)$ time, where d is the intrinsic dimension (e.g. doubling dimension), C a constant that depends on the curvature of $\mathcal{M}$, n is the number of points in $\mathcal{M}$, and D is the cost of computing the distance between a pair of points in $\mathcal{M}$. These data structures are especially useful for both finding near points to any given point, and for constructing nets of points at multiple resolutions, and they may be run in an online fashion.

A slower, simpler algorithm for constructing a net is to add points one a time if they are farther than δ from any point already in the net; finish when no more points can be added. For simplicity, this is the algorithm we have used in examples presented in this paper.

If a way of sampling points from $\mathcal{M}$ with respect to a measure μ with uniformly lower and upper bounded density with respect to the volume measure on $\mathcal{M}$, then it is easy to see that $O(\delta^{-d}\log(1/\delta))$ samples suffice, with high-probability, to obtain enough samples from which a δ-net may be extracted [33].

3.2 Dimension Reduction: Landmark Multidimensional Scaling (LMDS)

LMDS takes as input a set of landmarks $L \subset \mathcal{M}$ and a set of other points $Z \subset \mathcal{M}$, and constructs a map $\Phi : L \cup Z \to \mathbb{R}^d$ embedding L, Z into $\mathbb{R}^d$. LMDS computes all pairwise distances between points in L, and returns low dimensional coordinates which minimize the distortion given by

$$\sum_{l_i, l_j \in L} \left(\rho(l_i, l_j)^2 - ||\Phi(l_i) - \Phi(l_j)||^2_{\mathbb{R}^d}\right)^2 \tag{6}$$

over all possible mappings Φ. This is Multiscale Dimensional Scaling (MDS, [7]). At this point Φ is only defined on L: to extend it to $L \cup Z$, LMDS proceeds by computing the distances between each point in L and each point in Z, and for each point $z \in Z$ assigns coordinates $\Phi(z)$ which minimize

$$\sum_{l_i \in L} \left(\rho(l_i, z)^2 - ||\Phi(l_i) - \Phi(z)||^2_{\mathbb{R}^d}\right)^2 \tag{7}$$

over all possible choices d dimensional vectors $\Phi(z)$. For a full description of the algorithm, see [48]. If the distance ρ is Euclidean, the algorithm reduces to principal component analysis (PCA).

If the dimension d is unknown, one could learn d at this stage from observing the eigenvalues of the squared distance matrix obtained during MDS. Eigenvalues which are of order δ^2 correspond to directions along the manifold, and eigenvalues which are of order δ^4 or lower correspond to curvature (or noise). Thus, one could learn d by choosing a cutoff threshold depending upon δ (in fact this is done in example 5.5). For an extensive analysis on how to detect intrinsic dimensionality of data sets see [31, 2], and its use in the context of high-dimensional stochastic systems in view of global nonlinear dimension reduction and reaction coordinates see [45, 62].

3.2.1 Computational cost

The computational cost of this algorithm is $O((|L|^2 + |L| \cdot |Z|)D)$, where D is the cost of evaluating ρ at a pair of points. This cost comes directly from the number of distances computed. The point of LMDS compared to MDS (the case $|Z| = 0$) is that in the case of interest where $|L| \ll |Z|$, the cost is linear in $|L \cup Z|$ for LMDS instead of quadratic as in MDS [48]. The cost of computing d eigenvectors on a matrix of size $|L| \times |L|$ is $O(d|L|^2)$, which is negligible compared to $|L|^2 D$.

3.3 Least-squares switching maps

We will use the pseudoinverse (see [41]) to solve a least squares problem of finding the best linear transition map. If X and Y are each $2m \times d$ matrices (with chart images as rows, and with mean zero columns), then the matrix $T = X^\dagger Y$ minimizes $||XT - Y||_2$ over all $d \times d$ matrices.

Fix $y_k \in \Gamma$. In the construction algorithm that follows, for each connection $j \sim k$, we take a set of common landmarks $L_{k,j} = A_k \cup A_j$ and let $X = \Phi_k(L_{k,j})$ and $Y = \Phi_j(L_{k,j})$, with the d-dimensional vectors involved being the rows of these matrices. Since the mean of the rows of X (resp. Y) is not zero, we subtract from each row the mean $\mu_{k,j}$ (resp. $\mu_{j,k}$) of the rows of X (resp. Y). The charts C_k and C_j represent overlapping areas on $\mathcal{M}$, and so there will exist a matrix $T_{k,j} = X^\dagger Y$ which has small error. See again figure 1 for a detailed picture.

To simplify notation, we will combine the mean shifting and matrix multiplication into a single operator $S_{k,j}$

$$S_{k,j}(x) = (x - \mu_{k,j})T_{k,j} + \mu_{j,k} \tag{8}$$

To decide when to apply the switching maps, we will need to know when the simulator is in the region between two charts. To do this, we will calculate the distances to the *chart centers*, which we call $c_{k,j} = \Phi_k(y_j)$, i.e. the images of the net points via the low dimensional mappings.

3.3.1 Computational cost

The cost of computing the pseudoinverse is $O(2md^2)$ since we must compute the singular value decomposition of X, and $2m > d$. The cost of applying the switching map is $O(d^2)$.

3.4 Learning Phase

The first part of the ATLAS algorithm is the learning phase, in which we use the sample paths to learn local chart coordinates, local simulators and transition maps. In this part of the algorithm, we store all the

Learning Phase

$\widehat{\mathcal{S}} =$ **construction_phase**$(\{x_j\}, \rho, \mathcal{S})$
  $\{y_k\}_{k\in\Gamma} \leftarrow \delta-\text{net}(\{x_j\})$
  for $k \in \Gamma$
    % *create* $m+1$ *landmarks for LMDS* around y_k
    $\{a_{k,l}\}_{l=1..m} = \mathcal{S}(y_k, m, t_0)$
    $A_k = y_k \cup \{a_{k,l}\}$
  end
  for $k \in \Gamma$
    % *simulate* p *paths for estimating drift and diffusion coefficients around* y_k
    $\{x_{k,l}\}_{l=1..p} = \mathcal{S}(y_k, p, t_0)$
    $L_k = \bigcup_{i\sim k} A_i$
    $[L'_k, \{x'_{k,l}\}_{l=1..p}] = \text{LMDS}(L_k, \{x_{k,l}\}_{l=1..p}, \rho)$
    $\{\widehat{\mathcal{S}}.c_{k,j}\} \leftarrow \bigcup_{j\sim k} y_j$ in L'_k
    shift coordinates so $c_{k,k} = 0$
    $\widehat{\mathcal{S}}.\bar{b}_k \leftarrow \sum_l x'_{k,l}/pt_0$
    $\widehat{\mathcal{S}}.\bar{\sigma}_k \leftarrow (\text{Cov}(\{x'_{k,l}\})/t_0)^{1/2}$
    % *compute switching maps*
    for $j \sim k$, $j < k$
      $L_{k,j} = A_k \bigcup A_j$
      $L'_{k,j} = L_{k,j}$ in L'_k coordinates
      $L'_{j,k} = L_{j,k}$ in L'_j coordinates
      $\widehat{\mathcal{S}}.\mu_{k,j} \leftarrow \mathbb{E}[L'_{k,j}]$
      $\widehat{\mathcal{S}}.\mu_{j,k} \leftarrow \mathbb{E}[L'_{j,k}]$
      $\widehat{\mathcal{S}}.T_{k,j} \leftarrow (L'_{k,j} - \mu_{k,j})^\dagger (L'_{j,k} - \mu_{j,k})$
    end
  end

Figure 4: Main algorithm for constructing the ATLAS: it constructs the δ-net, computes chart embeddings, learns chart simulators from short sample paths, and transition maps.

information necessary for the global simulator in $\widehat{\mathcal{S}}$, and describe in the next section what it means to run this simulator. We will use the notation $\widehat{\mathcal{S}}.var$ to denote the variable var within the simulator $\widehat{\mathcal{S}}$. Recall that $\{x_j\}$ is a given, dense enough sample of $\mathcal{M}$ to produce a net at scale δ. Let $\mathcal{S}(y, p, t_0)$ denote running p paths of the simulator starting at y for time t_0. We treat all points on the charts (resulting from LMDS) as row vectors. We compute, as described in the algorithm in Figure 4, for each chart $k \in \Gamma$, a drift $\widehat{\mathcal{S}}.\bar{b}_k$ and a diffusion coefficient $\widehat{\mathcal{S}}.\bar{\sigma}_k$, effectively approximating the dynamics in the chart by that of a Itô diffusion with constant drift and diffusion coefficient in the chart space (i.e. after mapping to Euclidean space using the LMDS map), estimated from the p paths from $\mathcal{S}$ observed in the chart.

3.4.1 Computational cost

Each net point has order 2^d connections at most, by the properties of the δ-net. Let S denote the cost of running one simulation of length t_0. Thus, the computational cost of the construction phase, for each chart,

is of order

$$\underbrace{mS}_{\text{landmark simulation}} + \underbrace{pS}_{\text{path simulation}} + \underbrace{2^d mpD}_{\text{LMDS}} = (m+p)S + 2^d mpD \tag{9}$$

We note that the term $2^d mpD$ can be decreased to dpD since a d-dimensional plane may be estimated with only $O(d)$ points. Instead of using all cm points as landmarks, one could choose a (e.g. random) subset of these landmarks for the initial embedding (although all these landmarks will be needed later for computing T). All these steps are easily parallelized, so the per-chart cost above is also a per-processor cost if enough processors are available. Finally, observe that there at most $O(\delta^{-d})$ such charts (this follows from the property of the δ-net, which ensures that balls of radius $\delta/2$ centered at net points are disjoint).

3.5 The learned simulator

The second part of our construction is to actually define the ATLAS, i.e. the reduced simulator of the system, given the parameters learned in the first stage. In other words we must describe what a single step of time Δt looks like starting at a location x in chart i. Figure 5 contains the pseudocode for the algorithm implementing the strategy we now discuss (written assuming x is a row vector). Given the position at time t is x in chart i, the position at time $t + \Delta t$ is determined by first choosing i' so that $\widehat{\mathcal{S}}.c_{i',j}$ is closest to x, among all possible $\{\widehat{\mathcal{S}}.c_{k,j}\}_k$. If $i' \neq i$, then the coordinates of x (in chart i) are changed to coordinates in chart i' by applying the switching map $T_{i,i'}$. Now that x is in the coordinates of the chart i', a forward stochastic Euler step is taken using the drift and diffusion coefficients in chart i'. Finally, this Euler step is confined to the local chart by applying a "wall function" W; details will be discuss in the next section.

Learned Simulator

$(x, i') = \textbf{simulator_step}(x, i, \widehat{\mathcal{S}})$
  % *select new coordinate chart*
  $i' = \text{argmin}_j ||x - \widehat{\mathcal{S}}.c_{i,j}||^2_{\mathbb{R}^d}$
  if $i' \neq i$
    $x \leftarrow (x - \widehat{\mathcal{S}}.\mu_{i,i'})\widehat{\mathcal{S}}.T_{i,i'} + \widehat{\mathcal{S}}.\mu_{i',i}$
  end

  % *forward Euler step*
  $\eta \sim \mathcal{N}(0, I_d)$
  $x \leftarrow x + \widehat{\mathcal{S}}.\bar{b}_{i'}\Delta t + \eta\widehat{\mathcal{S}}.\bar{\sigma}_{i'}\sqrt{\Delta t}$

  % *prevent escape from local chart*
  if $|x| > 3\delta/2$
    $x \leftarrow W(x) := \frac{x}{|x|}\left(2\delta - \frac{\delta}{2}\exp\left(3 - \frac{2}{\delta}|x|\right)\right)$
  end

Figure 5: Algorithm for running the ATLAS, by combining local diffusions and linear transition map between charts.

3.5.1 Computational cost

The ATLAS runs in d dimensions, and does not require calls to the original simulator, so the running time now only depends on the local complexity of the homogenized problem. The number of simulation steps required to approach stationarity still depends upon the time it takes to converge to equilibrium, but so too did the original simulator.

If $c \approx 2^d$ is the maximum number of connections each net point has, the computational cost of each time step of the ATLAS is of order

$$\underbrace{d2^d}_{\text{distance computation}} + \underbrace{d^2}_{\text{forward step}} = (2^d + d)d\,. \tag{10}$$

4 Theoretical Results and Guarantees

In this section, we first introduce the minimum amount of material to precisely state the Theorem in section 4.1. Then we introduce the necessary mathematical objects to state the Lemmata used during the main proof in section 4.3. In section 4.4, we prove Theorem 4.1 and the Lemmata used.

4.1 Theorem Statement

Let $\{y_k\}_{k\in\Gamma}$ denote the set of net points. For each k we have the mapping Φ_k from $\mathcal{M}$ to $\mathbb{R}^d$ given by MDS. We will assume the projection to the tangent plane at y_k is invertible on a ball of radius 2δ, which implies Φ_k will also be invertible on a ball of radius 2δ (this is an intermediate step in the proof of Lemma 4.5). Let $\mathcal{A} = B_{2\delta}(0) \times \Gamma$, equipped with the transition maps $\{\Phi_j(\Phi_k^{-1})\}$, denote the atlas. We shift coordinates in each chart so that $c_{k,k} = \Phi_k(y_k) = 0$ (where $c_{k,j}$ was defined in section 3.3), so $\mathcal{A}$ contains all points in each chart within 2δ of $c_{k,k}$. Next we state some definitions in order to mathematically define a step of the ATLAS. If $i \sim j$, let $S_{i,j}(x) = (x - \mu_{i,j})T_{i,j} + \mu_{j,i}$ be the transition map between charts i and j defined in equation 8.

Let $W(x)$ be the wall function (which confines the simulator to $B_{2\delta}(0)$) defined by

$$W(x) = \begin{cases} x & |x| \le \frac{3\delta}{2} \\ \frac{2\delta x}{|x|} - \frac{\delta x}{2|x|}\exp\left(3 - \frac{2}{\delta}|x|\right) & |x| > \frac{3\delta}{2} \end{cases} \tag{11}$$

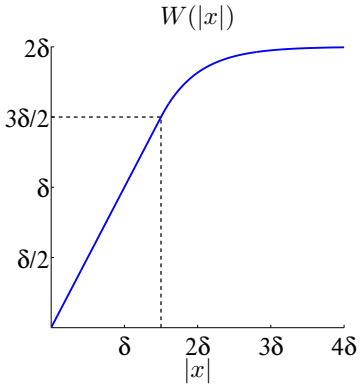

Figure 6: The wall function W.

There are other possible choices for W, but the main ingredients are: W is C^2, invertible, equal to the identity on a ball of radius $3\delta/2$, and takes $\mathbb{R}^d \to B_{2\delta}(0)$.

Let B_t be a standard Brownian motion in $\mathbb{R}^d$. The update rule for the ATLAS starting at (x_0, i_0) is

$$i_{k+1} = \text{argmin}_j \left|x_k - c_{i_k,j}\right| \tag{12}$$

$$x_{k+1} = W\left(S_{i_k,i_{k+1}}(x_k) + \bar{b}_{i_{k+1}}\Delta t + \bar{\sigma}_{i_{k+1}}B_{\Delta t}\right) \tag{13}$$

Define the ATLAS process $Z_k \in \mathcal{A}$ starting at $z_0 = (x_0, i_0)$ to be (x_k, i_k). See the algorithm in figure 5. We show in Lemma 4.2 that under the conditions of Theorem 4.1, Z_k has a unique stationary distribution μ on $\mathcal{A}$.

We will often wish to refer to the chart index associated with the following time step, starting from an $(x, i) \in \mathcal{A}$; we call this next chart index $i'(x, i)$, leaving off the (x, i) where it is obvious which initial condition is being talked about:

$$i'(x, i) = \text{argmin}_j \left|x - c_{i,j}\right| \tag{14}$$

Define for each i the continuous time process $\widehat{X}_t^x \in B_{2\delta}(0)$ by

$$\widehat{X}_t^x = W\big(S_{i,i'}(x) + \bar{b}_{i'}t + \bar{\sigma}_{i'}B_t\big) \tag{15}$$

We will often be referring to initial conditions of the form $z = (x, i) \in \mathcal{A}$, and for ease of notation we use $\widehat{X}_t^z$ to keep track of which $\widehat{X}$ process we are referring to, and what the starting location is. Keep in mind that $\widehat{X}_t$ lives in the local tangent plane.

Define a measure $\widehat{q}$ on $\mathcal{A}$ for each measurable set $E \subset \mathcal{A}$

$$\widehat{q}(E) = \int_{\mathcal{A}} \int_0^{\Delta t} \widehat{\mathbb{P}}\big[(\widehat{X}_t^z, i') \in E\big] dt \, \mathrm{d}\mu(z)\,, \tag{16}$$

where μ is the stationary distribution of Z_k on $\mathcal{M}$. By equation 15, one can see that $\widehat{q}$ is absolutely continuous w.r.t. Lebesgue measure (defined by the Lebesgue measures on each chart), and we will denote its Radon-Nikodym derivative, with some abuse of notation, $\widehat{q}(x)$. We show in Theorem 4.1 that $\widehat{q}$ is close to q, the stationary distribution of our original process Y_t.

Samples from $\widehat{q}$ may be generated by the ATLAS by running N steps of size Δt, and then one step of size $\delta t \sim$ Unif$(0, \Delta t)$. Samples of $\widehat{X}_{\delta t}$ are distributed according to $\int_0^{\delta t} \widehat{P}_t du$, as seen by integrating over the joint density, $\widehat{P}_u du$. Since Z is ergodic by lemma 4.2, $Z_N \to \mu$, thus there is an N for which our sampling plan (N steps plus a δt step) will approximate $\widehat{q}$ sufficiently well. Choosing N is not easy and depends on the problem, although this is a difficulty with the original simulator as well; in practice, one should choose N large enough that simulations reach a large fraction of the charts in the simulator. The reason for the random final time step δt (and the integral in equation 16) is to average over positions visited in between applications of the switching map.

Lastly, define a mapping back to the original space $G : \mathcal{A} \to \mathcal{M}$ by $G(x, i) = \Phi_i^{-1}(x)$ for each $(x, i) \in \mathcal{A}$. Define the push forward map G_* on $\widehat{q}$ by

$$G_*\widehat{q}(y) = \sum_{x:G(x)=y} \widehat{q}(x)dx \tag{17}$$

The set $\{x : G(x) = y\}$ is finite since the Φ_i's are one-to-one. This mapping G_* will allow us to compare q and $\widehat{q}$ as densities on $\mathcal{M}$.

Theorem 4.1. *Let $(Y_t)_{t\geq 0}$ be an Itô diffusion on a smooth compact connected d-dimensional manifold $\mathcal{M}$ with no boundary:*

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t\,, \tag{18}$$

with b, σ Lipschitz, and σ uniformly elliptic, i.e. there exists $\lambda > 0$ such that for all $x \in \mathcal{M}$ and $v \in T_x(\mathcal{M})$ $v^T\sigma(x)\sigma(x)^T v \geq \lambda|v|^2$. Let δ be small enough so that for every x the orthogonal projection $\mathcal{M} \to T_x(\mathcal{M})$ is invertible on a ball of radius 2δ on $T_x(\mathcal{M})$. Let q be density (with respect to volume measure on $\mathcal{M}$) of the unique stationary distribution of Y_t. Let $\widehat{q}$ be the density on $\mathcal{A}$ generated by the ATLAS, as defined above in (16). There exists constants c_1, c_2 such that if the number of sample paths satisfies $p > c_1(\tau, \mathcal{M})/\delta^4$ then with probability at least $1 - 2\exp(-\tau^2)$,

$$||q - G_*\widehat{q}||_{L^1(\mathcal{M})} \leq c_2\delta \ln(1/\delta) \tag{19}$$

Some remarks are in order:

(i) Lipschitz coefficients guarantee existence and uniqueness of solutions to (18); strong ellipticity (together with smoothness and connectivity of $\mathcal{M}$) guarantees that the process Y_t has a unique stationary distribution with smooth density (this latter assumption may be weakened to include hypo-elliptic systems).

(ii) From a computational perspective, we note that as G appears in the statement of the Theorem, it would be very useful to compute this mapping G. This is in general hard in arbitrary metric spaces, although one can always use the simple approximation $\widehat{G}(x,i) = y_i$ which approximates G at scale δ.

(iii) The dimension of the state space of the system does not appear in the Theorem - only the intrinsic dimension of $\mathcal{M}$ crucially affects the sampling requirements of the construction.

(iv) The ATLAS simulator is random, since it depends on the random paths collected to construct ATLAS, and so $\widehat{q}$ is random. The result states that with high probability $\widehat{q}$ is close to q in the sense of (19).

(v) In practice we are interested in applying our construction and the Theorem to processes Y_t that do not satisfy the assumptions above. This is possible when it is the case that above a certain timescale t_0, Y_t may be well approximated by a process $\widetilde{Y}_t$, representing Y_t *homogenized at time scale* t_0, which does satisfy the assumptions of the Theorem. In algorithmic terms, the microscale simulator (for Y_t) may not satisfy the conditions imposed above, but it may be well-approximated at timescales larger than t_0 by a simulator (for $\widetilde{Y}_t$) of the form (2) satisfying the conditions of the theorem on the timescale t_0. Our algorithm may then be applied to such a microscale simulator since it operates at time scales greater t_0 (its input is a set of paths of length t_0), and cannot therefore tell the difference between a microscale simulator for Y_t and one for its homogenization $\widetilde{Y}_t$. In this case the error in approximating the original process Y_t by $\widetilde{Y}_t$ may be added to the right hand side of (19).

We now turn to the proof of these results, after which we present many examples in section 5.

4.2 Algorithmic Complexity

Suppose that a single call to the original simulator of length t_0 costs S. The total number of points in the net Γ contains $O(\delta^{-d})$ points with constant depending on the volume of the manifold $\mathcal{M}$. From each point we will see that we must choose $p = O(\delta^{-4})$ to estimate the parameters of the ATLAS simulator to within accuracy δ. Assuming the expensive part of the construction algorithm is running the simulations, the total cost of construction is $O(S\delta^{-d-4})$.

During each call to the ATLAS, we will see from the proof that we must choose our timestep $\Delta t = t_0/\log(1/\delta)$. The $\log(1/\delta)$ term may be neglected, and we will do so in this discussion. At each call to the simulator, we must compute the distance to each neighbor, of which there are $O(2^d)$. Assuming each distance costs $O(d)$ flops, the total cost of running the ATLAS for time t_0 is $O(d2^d)$.

Comparing the running time of the original simulator S with the ATLAS amounts to comparing the cost of S to $d2^d$. The benefit of the ATLAS over the original one then clearly depends upon how expensive the original simulator was, which can depend on many factors: length of the timestep, cost of evaluating functions, ambient dimension, etc. One thing is clear - the cost of S depends on microscale properties, while $d2^d$ does not. Since the cost S varies for each problem, in our examples in section 5 we compare the ATLAS cost to the original simulator cost by comparing the size of the timestep alone.

The running time is the main benefit of the ATLAS, but not the only one. Another advantage is that long paths of the simulator can be stored using only d dimensions rather than the ambient dimension where $\mathcal{M}$ lives. Yet another benefit is that some postprocessing has already been done, for example consider the question "how long does Y_t spend near the state Y^*?". After running the original simulator one would have to compute a distance to Y^* for many data points. After running the ATLAS answering this question requires only computing distances from the few net points near Y^* to Y^* to obtain the result with accuracy 2δ, and the ATLAS already knows which parts of the paths are in charts near Y^*. Finally, ATLAS organizes the effective state space in terms of simple local low-dimensional models, enabling novel, interactive visualizations of the dynamical system, that can run on any (portable) device with a web browser, thanks to the compression achieved by the ATLAS: this is subject of ongoing work.

4.3 Preliminaries for Proofs

Before diving into the proof of Theorem 4.1, we need several definitions. In figure 35 we list a set of pointers to various definitions and notations used throughout. Let $X_t = \Phi_i(Y_t) \in B_{2\delta}(0)$. Let τ be the time when X_t first hits the boundary of $B_{2\delta}(0)$. Let $\Phi_{i,k}$ denote the k^{th} coordinate of Φ_i. Given that Y_t satisfies (18), a straightforward application of Itô's formula shows that $X_t = \Phi_i(Y_t)$ solves, for $t \leq \tau$, the Itô SDE

$$dX_t = b_i(X_t)dt + \sigma_i(X_t)dB_t \tag{20}$$

$$(b_i)_k(x) := \nabla\Phi_{i,k}(\Phi_i^{-1}(x)) \cdot b(\Phi_i^{-1}(x)) + \frac{1}{2}\sum_{j,l}\frac{\partial\Phi_{i,k}}{\partial x_j \partial x_l}(\Phi_i^{-1}(x))(\sigma\sigma^T)_{j,l}(\Phi_i^{-1}(x)) \tag{21}$$

$$(\sigma_i)_{k,j}(x) := \sum_l \frac{\partial\Phi_{i,k}}{\partial x_l}(\Phi_i^{-1}(x))\sigma_{l,j}(\Phi_i^{-1}(x)) \tag{22}$$

Let $\mathcal{L}$ be the generator (see [38]) for Y_t on $\mathcal{M}$, and, for $z = (x, i) \in \mathcal{A}$, $\mathcal{L}_z$ be the generator of X_t^z. $\mathcal{L}$ is uniformly elliptic on $\mathcal{M}$, since Y_t is uniformly elliptic on $\mathcal{M}$.

We will be comparing a number of simulators to bridge the gap between the simulation scheme $\widehat{\mathcal{S}}$ and the true simulator $\mathcal{S}$. We will do this by introducing intermediate processes in the local charts, and we will keep track of which chart the ATLAS process Z_k is in separately. For this reason we consider the following processes on $B_{2\delta}(0)$, started at $z = (x, i)$:

$$\bar{X}_t^z = \Phi_{i'}\left(\Phi_i^{-1}(x)\right) + \bar{b}_{i'}t + \bar{\sigma}_{i'}B_t \tag{23}$$

$$\widetilde{X}_t^z = S_{i,i'}(x) + \bar{b}_{i'}t + \bar{\sigma}_{i'}B_t \tag{24}$$

$$\widehat{X}_t^z = W(\widetilde{X}_t^z) \tag{25}$$

where i' is defined as in (14). The processes $\bar{X}_t, \widetilde{X}_t$ are natural stepping stones from X_t to $\widehat{X}_t$: $\bar{X}_t$ is the process which differs from X_t locally only in that it uses the learned drift and diffusion coefficients. $\widetilde{X}_t$ differs from $\bar{X}_t$ only in that it uses the learned transition map $S_{i,i'}$ rather than the true transition map $\Phi_{i'} \circ \Phi_i^{-1}$. Finally $\widehat{X}_t$ differs from $\widetilde{X}_t$ only because of the application of the local wall function W. Given any initial condition $z = (x, i)$, the three processes $\bar{X}_t^z, \widetilde{X}_t^z, \widehat{X}_t^z$ are solutions of SDE's with generators $\bar{\mathcal{L}}_z, \widetilde{\mathcal{L}}_z, \widehat{\mathcal{L}}_z$, respectively, on chart i'. These generators clearly depend on the chart i', but as we will see in Lemma 4.5, they will also depend on x, and we keep track of this by putting z as a subscript on the generators. We will prove that these generators are close to $\mathcal{L}_z$ for all $z \in \mathcal{A}$. Then we will show, using ideas from [34], that this is enough to imply bounds on the stationary distributions.

4.4 Proofs

Before we begin the proof of Theorem 4.1, we state some Lemmata which we will prove later, in order to keep the details until the end, while first showing the main ideas of the proof.

Lemma 4.2. *The process Z_k is ergodic with stationary distribution μ.*

Lemma 4.3. *For any smooth test function $F : \mathcal{A} \times B_{2\delta}(0) \to \mathbb{R}$ and initial condition Z_0,*

$$\frac{1}{n}\sum_{k=1}^{n}\int_0^{\Delta t} F(Z_k, \widehat{X}_t^{Z_k})dt \to \mathbb{E}_\mu\left[\mathbb{E}\left[\int_0^{\Delta t} F(z, \widehat{X}_t^z)dt\right]\right] \tag{26}$$

a.s. as $n \to \infty$. Here $\mathbb{E}_\mu$ means taking the expectation over the initial condition $z \sim \mu$, and $\mathbb{E}$ is taking the expectation over the transition probabilities of $\widehat{X}_t^z$.

The randomness in the statements of the following three Lemmata is that of the paths collected to learn the drift and diffusion coefficients used by $\bar{X}_t$ defined in (23).

Lemma 4.4. *There exists a constant C such that for any smooth test function $f : B_{2\delta}(0) \to \mathbb{R}$, and initial condition $z \in \mathcal{A}$, with probability at least $1 - 4\exp(-\tau^2)$*

$$\mathbb{E}\left[\frac{1}{\Delta t}\int_0^{\Delta t}(\mathcal{L}_z - \bar{\mathcal{L}}_z)f(\widehat{X}_t^z)dt\right] \leq C\delta\tau\sqrt{\ln(1/\delta)}||f||_{C^2} \tag{27}$$

Lemma 4.5. *There exists a constant C such that for any smooth test function f, and initial condition $z \in \mathcal{A}$, with probability at least $1 - 4\exp(-\tau^2)$*

$$\mathbb{E}\left[\frac{1}{\Delta t}\int_0^{\Delta t}(\bar{\mathcal{L}}_z - \widetilde{\mathcal{L}}_z)f(\widehat{X}_t^z)dt\right] \leq C\delta\ln(1/\delta)||f||_{C^2} \tag{28}$$

Lemma 4.6. *There exists a constant C such that for any smooth test function f, and initial condition $z \in \mathcal{A}$, with probability at least $1 - 4\exp(-\tau^2)$*

$$\mathbb{E}\left[\frac{1}{\Delta t}\int_0^{\Delta t}(\widehat{\mathcal{L}}_z - \widetilde{\mathcal{L}}_z)f(\widehat{X}_t^z)dt\right] \leq C\delta||f||_{C^2} \tag{29}$$

Proof of Theorem 4.1. This proof follows the ideas and techniques from [34] for proving large time convergence of numerical schemes, and we refer the reader to this reference for an overview of the previous substantial work in the area. Here we adapt part of the arguments contained in that paper to our setting. By assumption, the operator $\mathcal{L}$ is uniformly elliptic on $\mathcal{M}$ (for more details on how to define these operators on manifolds see [51], [26]). Let $\phi : \mathcal{M} \to \mathbb{R}$ be a smooth test function on $\mathcal{M}$ and define the average $\bar{\phi}$ by

$$\bar{\phi} = \int_{\mathcal{M}} \phi(y)q(y)dy \tag{30}$$

By construction $(\phi - \bar{\phi}) \perp \text{Null}(\mathcal{L}^*)$, and by the Fredhölm alternative there exists a unique solution ψ to the Poisson equation $\mathcal{L}\psi = \phi - \bar{\phi}$. Uniform ellipticity of $\mathcal{L}$ implies, via standard estimates [30], $||\psi||_{C^2} \leq C_{\mathcal{M},\lambda}||\phi||_\infty$.

For ease of notation, we will index everything by k, the step of the process $Z_k = (x_k, i_k) \in \mathcal{A}$. Let $\widehat{\mathcal{L}}_k = \widehat{\mathcal{L}}_{Z_{k-1}}$, $\widehat{X}_t^k = \widehat{X}_t^{Z_{k-1}}$, and $\psi_k = \psi \circ \Phi_{i_k}^{-1}$. Also let $\{B_t^k\}_{k=1}^\infty$ denote independent Brownian motions. The function ψ_k is smooth on $B_{2\delta}(0)$, so by Itô's formula:

$$\psi_k(\widehat{X}_{\Delta t}^k) - \psi_k(\widehat{X}_0^k) = \int_0^{\Delta t}\widehat{\mathcal{L}}_k\psi_k(\widehat{X}_t^k)dt + \int_0^{\Delta t}\nabla\psi_k(\widehat{X}_t^k)\widehat{\sigma}_{i_k}dB_t^k \tag{31}$$

By Itô's isometry, letting $||A(\cdot)||_{F,\infty} := ||\,||A(x)||_F\,||_{L^\infty(\mathcal{M})}$,

$$\mathbb{E}\left[\left(\int_0^{\Delta t}\nabla\psi_k(\widehat{X}_t^k)\widehat{\sigma}_{i_k}dB_t^k\right)^2\right] \leq \Delta t||\psi||_{C^1}^2||\widehat{\sigma}||_{F,\infty}^2 \tag{32}$$

Define the martingale M_n by

$$M_n = \frac{1}{n\Delta t}\sum_{k=1}^n\int_0^{\Delta t}\nabla\psi_k(\widehat{X}_t^k)\widehat{\sigma}_{i_k}dB_t^k \tag{33}$$

When calculating the variance of M_n, cross terms vanish by independence. Then from equation (32) we obtain the bound

$$\mathbb{E}[M_n^2] \leq \frac{1}{n\Delta t}||\psi||_{C^1}^2||\widehat{\sigma}||_{F,\infty}^2 \tag{34}$$

which implies $M_n \to 0$ a.s. as $n \to \infty$ by the martingale convergence theorem. Summing equation (31) and dividing by $n\Delta t$,

$$\frac{1}{n\Delta t}(\psi_k(Z_n) - \psi_k(Z_0)) = M_n + \frac{1}{n\Delta t}\sum_{k=1}^{n}\int_0^{\Delta t} \widehat{\mathcal{L}}_k\psi_k\big(\widehat{X}_t^k\big)dt \tag{35}$$

Since ψ is bounded, $\psi_k(Z_n)/n\Delta t \to 0$. Taking $n \to \infty$ on both sides and using Lemma 4.3,

$$0 = \mathbb{E}_\mu\left[\mathbb{E}\left[\frac{1}{\Delta t}\int_0^{\Delta t} \widehat{\mathcal{L}}_z\psi_z\big(\widehat{X}_t^z\big)dt\right]\right], \tag{36}$$

where, if $z = (x, i)$, $\psi_z = \psi \circ \Phi_{i'}^{-1}$ with i' defined as in equation (14). Using Lemma 4.4, 4.5 and 4.6, we have with probability at least $1 - 4\exp(-\tau^2)$,

$$\mathbb{E}_\mu\left[\mathbb{E}\left[\frac{1}{\Delta t}\int_0^{\Delta t} \mathcal{L}_z\psi_z\big(\widehat{X}_t^z\big)dt\right]\right] \le c\delta\tau\ln(1/\delta)||\phi||_\infty \tag{37}$$

Using the limit definition of the generator in [38], we can see that $\mathcal{L}_z\psi_z = \mathcal{L}\psi \circ \Phi_{i'}^{-1}$. Then since $\mathcal{L}\psi = \phi - \bar{\phi}$ and $\Phi_{i'}^{-1}(x) = G(z)$,

$$\left|\int_A \phi(G(z))\widehat{q}(z)dz - \int_{\mathcal{M}} \phi(y)q(y)dy\right| \le c\delta\tau\ln(1/\delta)||\phi||_\infty\,. \tag{38}$$

Since this holds for all smooth $\phi \in \mathbb{L}^\infty(\mathcal{M})$, we obtain

$$||q - G_*\widehat{q}||_{\mathbb{L}^1(\mathcal{M})} \le c\delta\tau\ln(1/\delta) \tag{39}$$

□

Proof of Lemma 4.2. Suppose first that we use the update rule starting at $Z_0 = (x_0, i_0)$

$$i_{k+1} = \text{argmin}_j\big|x_k + \eta u_k - c_{i_k,j}\big| \tag{40}$$

$$x_{k+1} = W\big(S_{i_k,i_{k+1}}(x_k) + \bar{b}_{i_{k+1}}\Delta t + \bar{\sigma}_{i_{k+1}}B_{\Delta t}\big) \tag{41}$$

$$Z_{k+1}^\eta = (x_{k+1}, i_{k+1}) \tag{42}$$

with $\{u_k\}$ being random variables drawn from Lebesgue measure on $B_{2\delta}(0)$, and $\eta > 0$. Note that $\eta = 0$ in the algorithm detailed in section 3 and in equation (12). If $\eta = 0$, the process Z_n^0 is not Feller continuous ($\mathbb{E}[Z_1^0]$ does not depend continuously on the initial conditions), a common assumption towards showing that Z_n has a stationary distribution. We start with fixing $\eta > 0$ and later will take $\eta \to 0$ and show that the η-dependent stationary measures converge to a new stationary measure. We temporarily suspend the use of the superscript η to simplify notation, till towards the end of the argument, when it will be relevant again when we let $\eta \to 0$.

First we show that the process Z_n is Feller continuous. Let f be a bounded function on $\mathcal{A}$ and $(x, i) \in \mathcal{A}$. Let $p_j(x, i)$ denote the probability of transitioning to chart j from i starting at x. Note that p_j is continuous and bounded provided $\eta > 0$. If $n = 1$,

$$\mathbb{E}_{(x,i)}[f(Z_1)] = \sum_{j\sim i} p_j(x, i)\mathbb{E}_{(x,i)}[f(Z_1)|i_1 = j] \tag{43}$$

Since $\widehat{X}_t^{Z_0}$ conditioned on i_1 is an Itô process, it is Feller continuous (see [38]). Thus $\mathbb{E}_{(x,i)}[f(\widehat{X}_{\Delta t}^{Z_0}, j)|i_1 = j]$ is continuous and bounded and therefore $\mathbb{E}_{(x,i)}[f(Z_1)]$ is continuous and bounded. By induction if $u(x, i) = \mathbb{E}_{(x,i)}[f(Z_n)]$ is continuous and bounded then

$$\mathbb{E}_{(x,i)}[f(Z_{n+1})] = \mathbb{E}_{(x,i)}\big[\mathbb{E}_{(y,j)}[f(Z_n)]\big|Z_1 = y, i_1 = j\big] = \mathbb{E}_{(x,i)}[u(Z_1)]\,, \tag{44}$$

which is continuous and bounded. Thus by induction on n, Z_n is Feller continuous for all $\eta > 0$.

Next we show the transition density of Z_n is tight for all n. Fix $\varepsilon > 0$. Let $z = Z_{n-1}$. Then $\widetilde{X}^z_{\Delta t}$ is Gaussian with mean $\widetilde{b}_z \Delta t$ and variance $\widetilde{\sigma}_z \widetilde{\sigma}_z^T \Delta t$ (see equation (69) for their definitions). $\sup_{z\in\mathcal{A}} \widetilde{b}_z$ and $\sup_{z\in\mathcal{A}} \widetilde{\sigma}_z$ are bounded. Thus there exists an R such that $\mathbb{P}[\widetilde{X}^z_{\Delta t} \in B_R(0)] > 1 - \varepsilon$ for all $z \in \mathcal{A}$, and thus $\mathbb{P}[Z_n \in W(B_R(0))] > 1 - \varepsilon$. $W(B_R(0))$ is compact which implies the transition density of Z_n is tight. The transition density is tight and Feller continuous, so by the Krylov-Bogolyubov Theorem [60] there exists an invariant measure.

Next we show that for any $(x, i), (y, j) \in \mathcal{A}$, if A_y is a neighborhood of y in chart C_j then $\mathbb{P}[Z_n \in A_y] > 0$ for large enough n. The $\delta-$net is a connected graph since $\mathcal{M}$ is connected. Thus there exists a finite length path $\{i_k\}$ with $i_k \sim i_{k-1}$, $i_0 = i$, $i_n = j$. The probability of such a path occurring is strictly positive since the probability of jumping from i_{k-1} to i_k is strictly positive for all k. The density of Z_n is strictly positive on the chart C_{i_n}, thus $\mathbb{P}[Z_n \in A_y] > 0$. Because the density of Z_n is positive on any open set in $\mathcal{A}$ for large enough n, the invariant measure is unique, and thus Z_n is ergodic.

Now let μ_η denote the stationary measure for Z^η_n for each $\eta > 0$. The family of measures μ_η is tight, and so there exists a subsequence $\{\mu_{\eta_k}\}_{k=1}^\infty$ which converges in probability to some measure μ (see [6]). It is left to show that μ is stationary for the process Z^0_n. Let $\varepsilon > 0$ and let f be a bounded function on $\mathcal{A}$. Let $\mu_k = \mu_{\eta_k}$. Let P_k denote the transition density for the process $Z^{\eta_k}_n$.

For ease of notation, for functions g, ν and transition kernels P,

$$g\nu = \int_{\mathcal{A}} g(x)\nu(x)dx \tag{45}$$

$$gP\nu = \int_{\mathcal{A}}\int_{\mathcal{A}} g(x)P(x, y)\nu(y)dxdy \tag{46}$$

Then $|f\mu_k - f\mu| \to 0$ by the bounded convergence theorem.

$$|fP_0\mu - f\mu| \le |fP_0\mu_k - fP_k\mu_k| + |fP_0\mu - fP_0\mu_k| + |f\mu_k - f\mu| \tag{47}$$

The last two terms go to zero as $k \to \infty$ because μ_k converges to μ in probability and f is bounded. It is left to show that $(P_0 - P_k)\mu_k \to 0$. Let E denote the boundary set defined by

$$E = \{(x, i) : \exists j \text{ with } |x| = |x - c_{i,j}|\}$$

The chart centers $c_{i,j}$ are a finite set and so E has μ measure zero. Let E_k denote the set E thickened by η_k:

$$E_k = \big\{(x, i) : \exists j \text{ with } \big||x| - |x - c_{i,j}|\big| < \eta_k\big\}$$

Fix $z = (x, i) \in \mathcal{A}$ and notice that the probability density starting at z, $P_k(z, \cdot)$, for any k is of the form

$$P_k(z, \cdot) = \sum_{j\sim i} p^k_j(z)\nu_j(\cdot)$$

with $p^k_j(z)$ being the probability of transitioning to chart j from i with $\eta = \eta_k$, and ν_j independent of k. For any j, ν_j is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}^d$, and $\nu_j(E_k) \to 0$. It follows that

$$\mu_k(E_k) \le \sup_{z\in\mathcal{A}} P_k(z, E_k) \to 0.$$

Since P and P_k agree on the set E^c_k, $|fP_k\mu_k - fP\mu| \to 0$. Thus, $fP\mu = f\mu$ for all test functions f, and therefore $P\mu = \mu$. □

Proof of Lemma 4.3. First let

$$I_n = \int_0^{\Delta t} F\big(Z_n, \widehat{X}^{Z_n}_t\big)dt \tag{48}$$

Then define a new Markov chain $Q_n = (Z_n, Z_{n+1}, I_n)$. Define a family of measures ν on $\mathbb{R}$ by

$$\nu(z_1, z_2, A) = \mathbb{P}[I_n \in A | Z_n = z_1, Z_{n+1} = z_2] \tag{49}$$

Let γ be a measure on $\mathcal{A} \times \mathcal{A} \times \mathbb{R}$ so that

$$\gamma(A) = \int_{\mathcal{A}\times\mathcal{A}\times\mathbb{R}} \mathbb{1}_A(r)\nu(z_1, z_2, dr)P(z_1, dz_2)\mu(dz_1) \tag{50}$$

Where P is the transition density for Z. Because Z_n is ergodic, $P^n(\delta_{(x,i)}, \cdot) \to \mu(\cdot)$ weakly. Then by the dominated convergence theorem as $n \to \infty$,

$$\int_{\mathcal{A}\times\mathcal{A}\times\mathbb{R}} \mathbb{1}_A(r)\nu(z_1, z_2, dr)P(z_1, dz_2)P^n(\delta_{(x,i)}, dz_1) \to \gamma(A) \tag{51}$$

The last statement shows that the density of Q_n converges weakly to γ, and so Q is ergodic. Pick $\phi(Q_n) = I_n$. Then by Birkhoff's ergodic theorem (see [58]),

$$\frac{1}{n}\sum_{k=1}^{n} \phi(Q_k) \to \int \phi d\gamma = \mathbb{E}_\mu \left[\mathbb{E}\left[\int_0^{\Delta t} F\big(z, \widehat{X}_t^z\big) dt \right] \right] \tag{52}$$

□

Proof of Lemma 4.4. Choose some $z = (x, i) \in \mathcal{A}$ and $f \in C^2$. Then the generators $\mathcal{L}_z, \bar{\mathcal{L}}_z$ are given by

$$\mathcal{L}_z f(y) = \sum_j \big(b_i(y)\big)_j \frac{\partial f}{\partial y_j}(y) + \frac{1}{2}\sum_j \sum_k \big(\sigma_i(y)\sigma_i^T(y)\big)_{j,k} \frac{\partial^2 f}{\partial y_j \partial y_k}(y)$$

$$\bar{\mathcal{L}}_z f(y) = \sum_j \big(\bar{b}_i\big)_j \frac{\partial f}{\partial y_j}(y) + \frac{1}{2}\sum_j \sum_k \big(\bar{\sigma}_i\bar{\sigma}_i^T\big)_{j,k} \frac{\partial^2 f}{\partial y_j \partial y_k}(y)$$

It suffices to show that $b_i(y)$ is close to $\bar{b}_i$ and $\sigma_i(y)\sigma_i^T(y)$ is close to $\bar{\sigma}_i\bar{\sigma}_i^T$ for each $y \in B_{2\delta}(0)$ and all i. Let $x_k = x'_{i,k}$ be random draws from $\Phi_i(Y_{\Delta t}^{y_i})$; These are samples of X_t starting at $c_{i,i} = 0$. Then as $p \to \infty$,

$$t\bar{b}_i \to \mathbb{E}\left[X_t\right] \qquad , \qquad t\bar{\sigma}_i\bar{\sigma}_i^T \to \mathrm{Cov}\left(X_t\right)$$

a.s. by the strong law of large numbers. Next in order to use finite sample bounds, we show that the random variables x_k are sub-gaussian with sub-gaussian norm $t_0\kappa(|\sigma|_{F,\infty} + t_0|b|_\infty)$ for some universal constant κ. To do this, we first show Y_{t_0} is sub-gaussian.

Rewrite the process Y_s by the definition of the Itô integral. Here we use a uniform partition of $(0, s)$ with n subintervals so $\Delta s = s/n$, $s_j = j\Delta s$, $Y_j = Y_{s_j}$ and z_j are independent standard random normal vectors in $\mathbb{R}^d$. Then $Y_s - Y_0$ can be written

$$Y_s - y_0 = \lim_{n\to\infty} \frac{1}{n}\sum_{j=0}^{n-1} b(Y_j)\Delta s + \sigma(Y_j)\sqrt{\Delta s} z_j \tag{53}$$

Note that we can always think of equations (18),(53) as being in $\mathbb{R}^D$ by the Whitney Embedding Theorem. If one is concerned about how to make sense of equation (53) on a manifold in $\mathbb{R}^D$, see [26]. Using proposition 5.10 in [55] on the right hand side of equation (53), we can see that there is a universal constant c so that for each n,

$$\mathbb{P}\left[\left| \frac{1}{n}\sum_{j=0}^{n-1} b(Y_j)\Delta s + \sigma(Y_j)\sqrt{\Delta s} z_j \right| > \alpha \right] \leq \exp\left(\frac{-c\alpha^2}{s(|\sigma|_{F,\infty} + s|b|_\infty)} \right) \tag{54}$$

Taking $n \to \infty$ we conclude that the sub-gaussian norm of Y_{t_0} is bounded by $\sqrt{t_0\kappa(|\sigma|_{F,\infty} + t_0|b|_\infty)}$ for a universal constant κ. Then X_{t_0} is also sub-gaussian with at most the same sub-gaussian norm since Φ_i is a projection. Then $|t_0\bar{b}_i - t_0\mathbb{E}[\bar{b}_i]|$ can be written as a sum of mean zero sub-gaussians and by [55] there exists a c_1 such that,

$$\mathbb{P}[\,|\bar{b}_i - \mathbb{E}[\bar{b}_i]| < \varepsilon] \geq 1 - e \cdot \exp\left(-c_1\varepsilon^2 p t_0\right) \tag{55}$$

Again by [55] bounds on finite sample covariance estimation yields for some c_2,

$$\mathbb{P}\Big[\,\big|\big|\bar{\sigma}_i\bar{\sigma}_i^T - \mathbb{E}\left[\bar{\sigma}_i\bar{\sigma}_i^T\right]\big|\big|_2 < \varepsilon \big|\big|\mathbb{E}\left[\bar{\sigma}_i\bar{\sigma}_i^T\right]\big|\big|_2\Big] \geq 1 - 2e^{-c_2 d\alpha^2} \tag{56}$$

provided $p > d\alpha^2/\varepsilon^2$. Notice that t_0 appears in the bound in equation (55), but not in equation (56). This is due to the fact that for $t_0 \ll 1$, the mean is much smaller than the standard deviation (and thus harder to estimate). Estimating the covariance to within accuracy ε takes $O(d/\varepsilon^2)$ samples, but estimating the drift to within accuracy ε takes $O(1/t_0\varepsilon^2)$ samples. Assuming $t_0 = \delta^2 \ll 1/d$, the mean will be more difficult to estimate. For simplicity we will assume that the covariance has the same bound as the drift.

Next we must ensure that the probabilistic bound holds for all indices $i \in \Gamma$. Since the volume of $\mathcal{M}$ is fixed, $|\Gamma| = c_3(1/\delta)^d$ for some c_3. Next set the accuracy to $\varepsilon = \delta\ln(1/\delta)$, and the confidence $\tau^2 = c_1\varepsilon^2 p t_0 - (1 + \ln(c_3) + d\ln(1/d))$. When we take a union bound over $i \in \Gamma$ we have with probability at least $1 - 2e^{-\tau^2}$,

$$|\bar{b}_i - \mathbb{E}[\bar{b}_i]| < \delta\ln(1/\delta) \text{ and } \big|\big|\bar{\sigma}_i\bar{\sigma}_i^T - \mathbb{E}\left[\bar{\sigma}_i\bar{\sigma}_i^T\right]\big|\big|_2 < \delta\ln(1/\delta) \tag{57}$$

$$\text{if } p > \frac{c_1}{t_0\delta^2}\left(\frac{\tau^2 + 1 + \ln(c_3)}{\ln(1/\delta)^2} + \frac{d}{\ln(1/\delta)}\right) \tag{58}$$

We can think of equation (58) as telling us $p > c_4/\delta^4$ (up to $\ln(1/\delta)$ factors) since $t_0 = \delta^2$ and $\tau, d, \ln(1/\delta)$ all behave like $O(1)$ constants.

Since Φ_i is smooth, b_i, σ_i are Lipschitz and bounded because b, σ are Lipschitz and bounded by some constant M. By the Cauchy-Schwarz inequality and Itô's isometry,

$$\mathbb{E}[|X_{t_0}|^2] \leq M^2 t_0 + O({t_0}^{3/2}) \qquad , \qquad \mathbb{E}\left[\int_0^{t_0} |X_s|^2 ds\right] \leq \frac{1}{2}M^2{t_0}^2 + O({t_0}^{5/2}) \tag{59}$$

Let $A = \int_0^{t_0} b_i(X_s) - b_i(0)ds$ and $B = \int_0^{t_0} \sigma_i(X_s)dB_s$. Then by Jensen's inequality,

$$|\mathbb{E}[X_{t_0}] - t_0 b(0)|^2 \leq \mathbb{E}[|A|^2] \leq \mathbb{E}\left[t_0\int_0^{t_0} |b_i(X_s) - b_i(0)|^2 ds\right] \leq \frac{1}{2}C^2M^2{t_0}^3 + O({t_0}^{7/2}) \tag{60}$$

Dividing by t_0 and taking a square root,

$$|\mathbb{E}[\bar{b}_i] - b_i(0)| \leq \sqrt{\frac{t_0}{2}}MC + O({t_0}^{3/4}) \tag{61}$$

By Itô's isometry we have

$$\mathbb{E}[|B|^2] \leq M^2 t_0 \tag{62}$$

Combining equations (62) and (60),

$$|\mathrm{Cov}(A+B) - \mathrm{Cov}(B)|_2 \leq \mathbb{E}[|A|^2] + 2\mathbb{E}[|A|^2]^{1/2}\mathbb{E}[|B|^2]^{1/2} \leq \sqrt{2}CM^2{t_0}^2 + O({t_0}^{9/4}) \tag{63}$$

Using Itô isometry and the Lipschitz condition on σ_i,

$$|\mathrm{Cov}(B) - t_0\sigma_i(0)\sigma_i^T(0)|_2 = \left|\mathbb{E}\left[\left(\int_0^{t_0}\sigma_i(X_s) - \sigma_i(0)dB_s\right)\left(\int_0^{t_0}\sigma_i(X_s)dB_s\right)^T\right.\right. \tag{64}$$

$$\left.\left. + \left(\int_0^{t_0}\sigma_i(0)dB_s\right)\left(\int_0^{t_0}\sigma_i(X_s) - \sigma_i(0)dB_s\right)^T\right]\right| \tag{65}$$

$$\leq \sqrt{2}KM{t_0}^{3/2} + O({t_0}^{7/4}) \tag{66}$$

Combining equations (61) with the concentration inequality (57) along with $t_0 = \delta^2$ implies for some c_5 with probability at least $1 - 2e^{-\tau^2}$,

$$|\bar{b}_i - b_i(0)| \leq c_5\delta\ln(1/\delta) \tag{67}$$

Finally, combining equations (63), (66) with the concentration inequality (57) yields that for some constant c_6 with probability $1 - 2e^{-\tau^2}$,

$$\left|\left|\bar{\sigma}_i\bar{\sigma}_i^T - \sigma_i(0)\sigma_i^T(0)\right|\right|_2 \leq c_6\delta\ln(1/\delta) \tag{68}$$

The Lipschitz conditions on b_i and σ_i yield the result.

□

Proof of Lemma 4.5. Fix a starting location $z = (x, i) \in \mathcal{A}$. We can write an SDE for $\widetilde{X}_t^z$ starting at $\Phi_{i'}\left(\Phi_i^{-1}(x)\right)$ in the next chart i' by

$$\begin{aligned} d\widetilde{X}_t^z &= \widetilde{b}_z dt + \widetilde{\sigma}_z dB_t \\ \widetilde{b}_z &= \frac{1}{\Delta t}\left(S_{i,i'}(x) - \Phi_{i'}\left(\Phi_i^{-1}(x)\right)\right) + \bar{b}_i \\ \widetilde{\sigma}_z &= \bar{\sigma}_{i'} \end{aligned} \tag{69}$$

Writing this equation in this form spreads the transition error out over the course of one timestep of length Δt. Thus, proving $\widetilde{\mathcal{L}}_z$ is close to $\bar{\mathcal{L}}_z$ reduces to showing that the transition error is sufficiently small after dividing by Δt (so that it can be combined in the drift term). Allowing the transition error to affect the drift forces us to have the drift $\widetilde{b}$ (and thus $\widetilde{\mathcal{L}}$) depend on the starting location z.

By the Whitney embedding theorem, $\mathcal{M}$ can be smoothly embedded into $\mathbb{R}^D$ for $D \geq 2d$. In $\mathbb{R}^D$, the LMDS mapping Φ_i reduces to principal component analysis, which is simply a projection onto the top d eigenvectors of the covariance matrix of the local landmarks. Thus we can think of Φ_i as a matrix acting on vectors. To be consistent with the algorithm, vectors will be written as row vectors and the matrix Φ_i will act on the left.

Fix $k \in \Gamma$. Let $\Pi_k \in \mathbb{R}^{D\times d}$ denote the projection from $\mathcal{M}$ onto T_{y_k}, the tangent plane of $\mathcal{M}$ at y_k. Then Π_k is invertible on a ball of radius 2δ on T_{y_k} by assumption. Also since $\mathcal{M}$ is smooth, Taylor's theorem tells us that for some c_1 and all $x \in \mathcal{M}$ near y_k,

$$|x\Pi_k - x| \leq c_1|x - y_k|^2 \tag{70}$$

Let $L_k = \{l_{k,i}\}$ denote the collection of landmarks associated with the neighbors of y_k, and μ_k denote their mean. The matrix Φ_k minimizes the squared error on the landmarks given by:

$$\sum_i |l_{k,i}\Phi_k - l_{k,i}|^2 \tag{71}$$

Inserting Π_k in place of Φ_k into equation (71) yields a bound of $c^2\delta^4$. The landmarks are well spread through the space by construction and the ellipticity condition. Then with high probability (in fact, at least

$1 - e^{-\tau^2}$ by [56], since $p \gtrsim d \log d$) $\widetilde{L} = (L - \mu)\Pi = \{\widetilde{l}_i\}$ must have smallest singular value at least δ, and thus any vector v in the tangent plane can be written $v = a\widetilde{L}$ with $a = v\widetilde{L}^\dagger$. The bound on the singular values implies $|a| \leq \delta^{-1}|v|$. Then using Cauchy-Schwarz,

$$|v\Phi_k - v| \leq |v| c_1 \delta \,, \tag{72}$$

which implies, since Φ_k, Π_k are projections, $||\Phi_k - \Pi_k||_2 \leq c_1\delta$. Let $j \in \Gamma$ such that $j \sim k$. By a Taylor expansion, $||\Pi_k - \Pi_j|| \leq c_2\delta$ for some c_2. Thus there exists a constant c_3 such that

$$||\Phi_k - \Phi_j||_2 \leq c_3\delta \tag{73}$$

The properties (72) (73) allow us to treat Φ_k like Π_k, the projection onto the tangent plane. Also since $||\Phi_k - \Pi_k||_2 \leq c_1\delta$ and $\delta \ll 1$, Φ_k will be invertible whenever Π_k is since Π_k has singular values equal to 1.

Next let $A = A_{k,j} = \{a_{k,i}\} \cup \{a_{j,i}\}$ be the collection of landmarks common to L_k and L_j. Let $\mu = \mu_{k,j}$ be the mean of A. Now we can write $T_{i,j}$ as

$$T_{i,j} = [(A - \mu)\Phi_i]^\dagger (A - \mu)\Phi_j \tag{74}$$

By definition of the pseudoinverse, $T_{i,j}$ minimizes

$$||(A - \mu)\Phi_i T - (A - \mu)\Phi_j||_2 \tag{75}$$

over all choices of $T \in \mathbb{R}^{d\times d}$. Choose T to be the restriction of Φ_j onto chart C_i. Then $T \in \mathbb{R}^d \times \mathbb{R}^d$ and

$$||(A - \mu)\Phi_i T - (A - \mu)\Phi_j||_2 = ||((A - \mu)\Phi_i - (A - \mu))(\Phi_j - \Phi_i)|| \leq c_1 c_2 \delta^3 \tag{76}$$

Since T is a possible choice for $T_{i,j}$,

$$||(A - \mu)\Phi_i T_{i,j} - (A - \mu)\Phi_j||_2 \leq c_1 c_3 \delta^3 \tag{77}$$

The matrix of landmarks $(A - \mu)$ spans the chart C_i, so there is a constant c_4 such that for any x in the chart C_i,

$$|S_{i,j}(x) - \Phi_j(\Phi_i^{-1}(x))| \leq c_4\delta^3 \tag{78}$$

Using $\Delta t = \delta/\ln(1/\delta)$ and equation (78), the result follows. □

Proof of Lemma 4.6. Fix a starting location $z = (x, i) \in \mathcal{A}$. Then the process $\widetilde{X}_t^z$ is the solution of an SDE on chart i' with smooth coefficients. Thus, $\widehat{X}_t^z = W\big(\widetilde{X}_t^z\big)$ is also the solution of an SDE on chart i' with smooth coefficients:

$$d\widehat{X}_t = \widehat{b}(z, \widehat{X}_t^z)dt + \widehat{\sigma}(z, \widehat{X}_t^z)dB_t \tag{79}$$

Using Itô's formula on $W(\widetilde{X})$,

$$\widehat{b}_j(z, \widehat{X}_t^z) = \sum_k \frac{\partial W_j}{\partial x_k}(\widetilde{X}_t^z)\widetilde{b}_k(z) + \frac{1}{2}\sum_k\sum_l \frac{\partial^2 W_j}{\partial x_k \partial x_l}(\widetilde{X}_t^z)(\widetilde{\sigma}\widetilde{\sigma}^T)_{k,l}(z) \tag{80}$$

$$\widehat{\sigma}_{j,l}(z, \widehat{X}_t^z) = \sum_k \frac{\partial W_j}{\partial x_k}\widetilde{\sigma}_{k,l}(z) \tag{81}$$

with $\widetilde{b}(z) = \widetilde{b}_z, \widetilde{\sigma}(z) = \widetilde{\sigma}_z$ defined as in the proof of Lemma 4.5. Note that since W is invertible, we could replace $\widetilde{X}_t^z$ with $W^{-1}(\widehat{X}_t^z)$ so that $\widehat{b}, \widehat{\sigma}$ can be thought of as a function of z and $\widehat{X}_t^z$. Direct computation shows that for some c_1, c_2,

$$\sum_k \left(\frac{\partial W_j}{\partial x_k}(\widetilde{X}_t)\right)^2 \leq c_1 \qquad , \qquad \sum_k\sum_l \left(\frac{\partial^2 W_j}{\partial x_k \partial x_l}(\widetilde{X}_t)\right)^2 \leq \frac{c_2}{\delta^2}$$

Let E_t denote the set $\left\{t : |\widetilde{X}_t^z| > \frac{3\delta}{2}\right\}$. By definition of W, $\widehat{b}$ and $\widetilde{b}$ agree on E_t^c. The index i' is chosen so that $|x - c_{i,i'}| < \delta$. As the switching map $S_{i,i'}$ makes error $O(\delta^3)$ by equation 78, $|\widetilde{X}_0^z| \leq \delta + O(\delta^3)$. In order that $|\widetilde{X}_t^z| > 3\delta/2$, the Brownian motion must push the process at least $O(\delta)$ in time t. In other words there are constants c_3, c_4 such that,

$$\mathbb{P}[E_t] \leq \mathbb{P}[|B_t| > c_3\delta] \leq \exp(-c_4\delta^2/4t) \tag{82}$$

Next we can bound the effect of the boundary function W on the drift and diffusion terms:

$$\mathbb{E}\left[\int_0^{\Delta t} \left|\widehat{b} - \widetilde{b}\right|^2 (z, \widehat{X}_t^z)dt\right] = \mathbb{E}\left[\int_0^{\Delta t} \mathbb{1}_{E_t} \left|\widehat{b} - \widetilde{b}\right|^2 (z, \widehat{X}_t^z)dt\right] \leq \frac{c_5\Delta t}{\delta^2} \exp\left(-c_4\frac{\delta^2}{\Delta t}\right)$$

for some new constant c_5. By equations (81), (82) and the fact that $\widehat{\sigma}$ agrees with $\widetilde{\sigma}$ on E_t^c,

$$\mathbb{E}\left[\int_0^{\Delta t} ||\widehat{\sigma}\widehat{\sigma}^T - \widetilde{\sigma}\widetilde{\sigma}^T||_F^2 (z, \widehat{X}_t^z)dt\right] \leq \Delta t c_6 \exp\left(-c_4\frac{\delta^2}{\Delta t}\right) \tag{83}$$

The result follows for $\Delta t = \delta/\ln(1/\delta)$. □

5 Examples

5.1 Simulator Comparison

5.1.1 Multiscale Transition Probability Comparisons

In order to see how well the ATLAS works we will need to have a criterion for comparing simulators. Since we are interested in the behavior of the system over multiple timescales, we will simulate 10,000 paths from each simulator and record the positions at times $\{t_k := 2^k\}$. The smallest time scale will be at the size of one step of the original simulator and the largest time scale will be at some time T (example dependent) at which point systems have reached equilibrium.

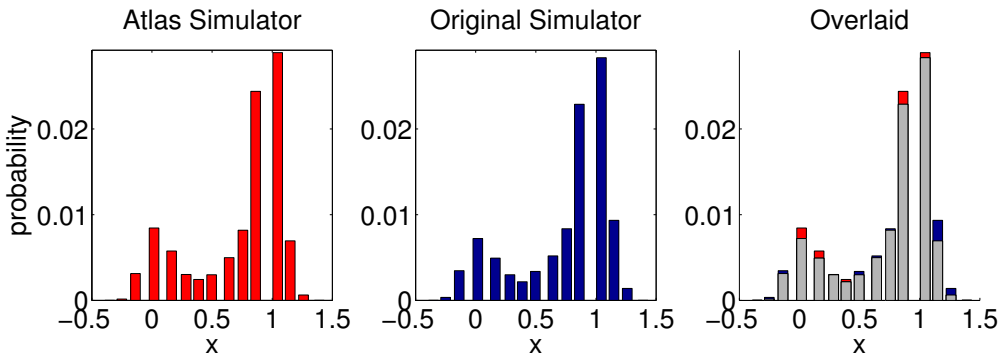

Figure 7: Comparing distributions obtained from two simulators at time $T = 0.2$ (orginal and ATLAS) in example 5.2.2.

In order to understand motivation for how to compare simulators, we start with a 1-d example. For a fixed $t_j \leq T$, we can bin samples into equal spaced bins, and compare them. Next we would like to compare the probabilities of landing in each bin as in figure 7 by overlaying them. We can next vary k (and thus t_k) to obtain overlaid bar graphs for multiple time scales as in figure 8. We see that the real quantity of interest

is the difference between these two histograms, and we will sum the absolute values of their difference to approximate the L^1 distance between the measures these histograms represent.

Our next goal is to generalize this to high dimensional spaces. Here the bins we use can be given naturally by the ATLAS we construct. Instead of using a "hard" binning procedure by assigning each point to the closest bin, we will assign smooth weights to the nearest neighbors. This smooth binning procedure will help to wash out the small scale errors we make, so that we can measure the large scale errors.

The first step is to explain the smooth map which takes a distribution ν on $\{x_i\}_{i=1}^n$ to a distribution μ on a set $\{y_j\}$. We think of $\{y_j\}$ as a coarser binning of the distribution ν on $\{x_i\}$. First assign weights $w_{i,j}$ to each (x_i, y_j) pair given by:

$$w_{ij} = \begin{cases} \exp\left(\frac{-|x_i - y_j|^2}{\delta^2}\right) & |y_i - x_j| < 2\delta \\ \\ 0 & \text{otherwise} \end{cases}$$

Then we normalize the weights so that they sum to 1 when summed over j.

$$\mu_j = \sum_i \nu_i \frac{w_{i,j}}{\sum_j w_{i,j}} \tag{84}$$

Fix a time slice t_k, then assign equal weights $\nu_i = 1/n$ to the set of samples $\{x_i\}_{i=1}^n$ given by the original simulator, and map them to a distribution μ on the net Γ using (5.1.1),(84) and the distance function in the ambient space. Next we will assign equal weights to the samples $\{\widehat{x}_i\}_{i=1}^n$ from the ATLAS and map them to weights $\widehat{\mu}$ on the net Γ using the euclidean chart distances.

Once we have μ, $\widehat{\mu}$, we could compare them directly. However, we know that the ATLAS makes errors on this spatial scale, and so we would like to smooth these distributions out to a coarser net with $\delta_c \geq \delta$. This will also allow us to compare simulators with varying δ while keeping the number of bins fixed. For each example, we will fix a coarse grained δ_c equal to the largest δ used for that example, and obtain a net $\{z_l\}$. Then we can push μ, $\widehat{\mu}$ to distributions p, $\widehat{p}$ on $\{z_l\}$ again using (5.1.1),(84) and the distance function in the ambient space. This gives us two probability distributions, one for each simulator, at time t_k on the coarse net.

Given a single initial condition, we will calculate the L^1 distance between p and $\widehat{p}$ for each time slice t_k. Then we will repeat this procedure for 10 fixed initial conditions (randomly chosen) to compare the transition densities over a wide range of time scales and initial conditions. In examples where only one ATLAS is used, we plot one thin colored line for each initial condition, then a thick line representing the mean $\pm$ one standard deviation (see figures 14, 20, 25, 31). In examples where we compare many ATLAS

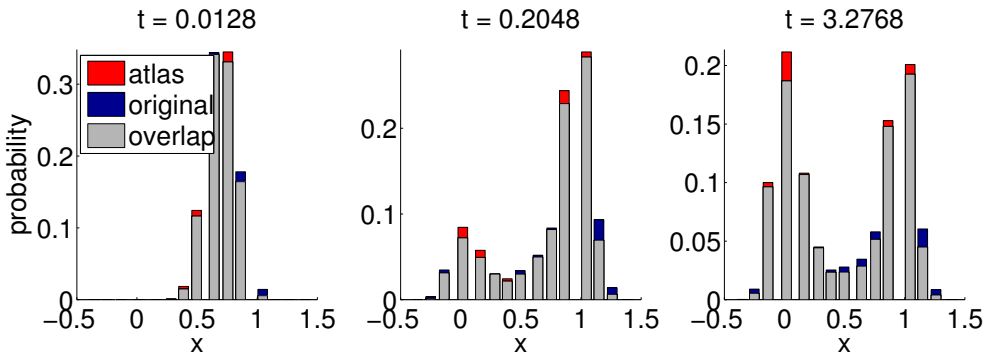

Figure 8: Comparing overlaid distributions obtained from two simulators at some multiscale times 2^k (originally and ATLAS) in example 5.2.2.

simulators, we plot only the thick line representing the mean ± one standard deviation (see figures 11, 18, 23, 29).

5.1.2 Transition Time Comparisons

Another quantity of interest in many stochastic dynamical systems are (expected) transition times between metastable states. Metastable states are subsets of the state space where the system spends a significant amount of time before escaping. An example includes a region near at the bottom of either well of the potential function in the example considered in section 2.3.2, and discussed in detail below. These transitions between a metastable state (set) and another are often one of the most important characteristics of large-time dynamics in the systems we consider. The expected time between these transitions, i.e. the average time spent in a metastable state before jumping to another one - is a fundamental statistic of the system, and is a function of two given metastable states (sets).

Therefore another way to compare the ATLAS with the original simulator is to compare these expected transition times. We start by running 12 extremely long paths (100 times longer than those run in 5.1.1) from both simulators; these paths have on the order of thousands of transitions. Next we identify regions of interest, which requires knowledge of the problem ahead of time. Once this is done we can classify points in the original state space as belonging to region 1, region 2, region 3 (if there are three states), or none. Then all points in the long simulations are classified. Last, we scan through the list and calculate the transition times. Calculating transition times is best explained through an example. Suppose our simulation now looks like:

$$0, 1, 1, 0, 0, 1, 0, 2, 0, 3, 3, 0, 1 \tag{85}$$

Denote a transition time from region i to region j by $\tau_{i,j}$. Start by skipping to the first time the simulator enters a region (region 1 in this case). Then it takes 6 timesteps to reach another region, region 2. Count this as a sample of $\tau_{1,2}$ equal to 6 timesteps. Next it takes 2 steps to go from region 2 to region 3, so count this as a sample of $\tau_{2,3}$ equal to 2 timesteps. Next we are in region 3, and it takes 3 steps before reaching region 1, so count this as a sample of $\tau_{3,1}$ equal to 3 timesteps. There are no samples of $\tau_{1,3}$ or $\tau_{3,2}$. The long paths we calculate will have many such samples, and we can average the value of these samples, and also average over the 12 paths we have run. We can then plot estimated values for $\mathbb{E}[\tau_{i,j}]$, the expected transition times for the original simulator, and estimated values of $\mathbb{E}[\widehat{\tau}_{i,j}]$, the expected transition times for the ATLAS.

5.2 One Dimensional Example

5.2.1 Smooth Potential

The first example presented is a simple one dimensional two-well example. We will use the potential

$$U_1(x) = 16x^2(x-1)^2$$

and use a simulator which approximates

$$dX_t = -\nabla U_1(X_t)dt + dB_t$$

using an Euler-Maruyama scheme which takes timesteps of size 0.005. A sample path of this system is shown in figure 9. The initial point set we use to generate the δ-net is linearly spaced points with spacing 0.01. It is important to note that the distribution of the initial point set does not play an important role in the resulting ATLAS. The ATLAS algorithm performs equally well on any initial point set that has no holes of size order δ. We subsample this initial point set to obtain a δ-net with δ parameter 0.1 using the brute force method described in section 3.1.

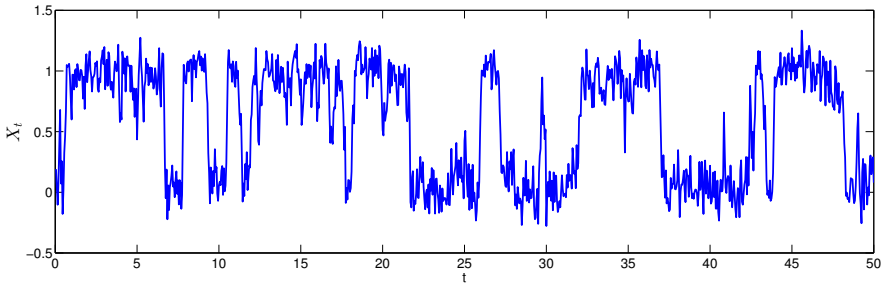

Figure 9: Sample trajectory of X_t for the two well example 5.2.1.

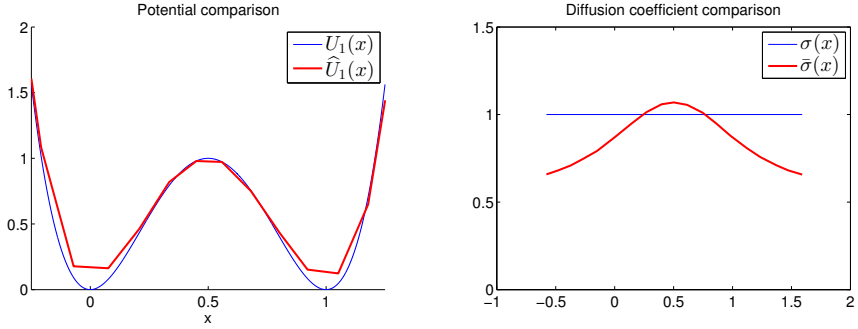

Figure 10: Left: original potential U (shown in blue) and effective potential of the ATLAS $\widehat{U}$ (shown in red). Right: comparing original diffusion coefficient (blue) with that of the ATLAS(red) with $\delta = 0.1$ in example 5.2.1.

Once we run the ATLAS algorithm in this case, it is simple to map estimated drift vectors from the chart coordinates back to the original space. In general for an arbitrary metric space this is a hard problem, but in 1-d we need only multiply by ± 1 to undo MDS. In 1-d, the estimated drift vectors can easily be integrated to obtain an effective potential $\widehat{U}$ for the system. We can also bring back the diffusion coefficients and see how they compare to the truth. Inverting MDS and comparing the coefficients we obtain with the true coefficients of the underlying system is a procedure we will only be able to do for this 1-d system, but it gives interesting insight into the working of the homogenizing nature of the ATLAS.

See Figure 10 showing the resulting comparisons between drift and diffusion coefficients.

Next we generate four nets (and four ATLAS simulators) with δ values 0.05, 0.10, 0.15 and 0.20 by subsampling from a fine mesh. In each example we have used $p = 10,000$ simulations per net point, and $t_0 = \delta^2$. The number of landmarks is irrelevant because as long as $m \geq 1$, there will be enough landmarks to exactly recover the local space. When simulating, we set the simulation time step $\Delta t = \delta^2/5$. Then for each of 10 randomly chosen staring locations, we run $10,000$ long paths up to time $T = 50$. Using the simulator comparison method from section 5.1, we obtain figure 11. As we expect from theorem 4.1, the long time error is decreasing as δ decreases. Figure 11 shows that the transition kernels are close for all time scales, which is a stronger experimental result than given by theorem 4.1. Theorem 4.1 only tells us that the stationary distributions are $O(\delta \log(1/\delta))$ far from each other. We can also compare the rates directly as seen in figure 12.

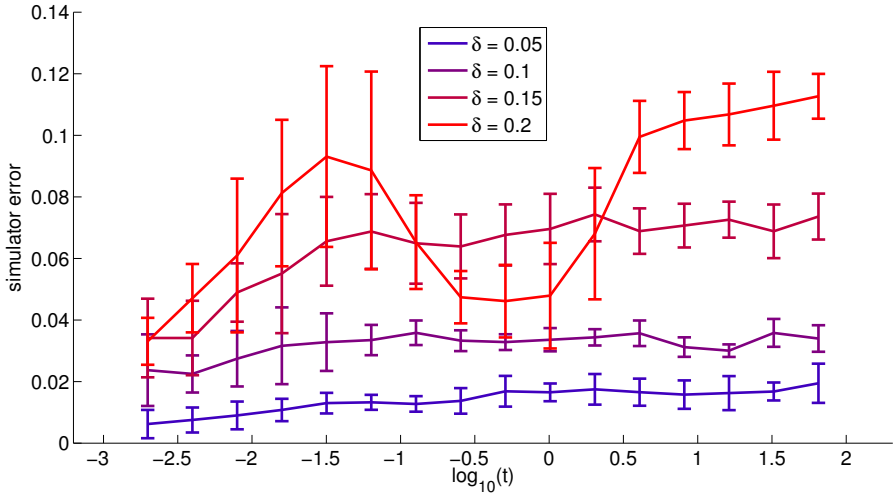

Figure 11: Simulator comparison for example 5.2.1. Each line represents the average simulator error for a single net of the specified δ value.

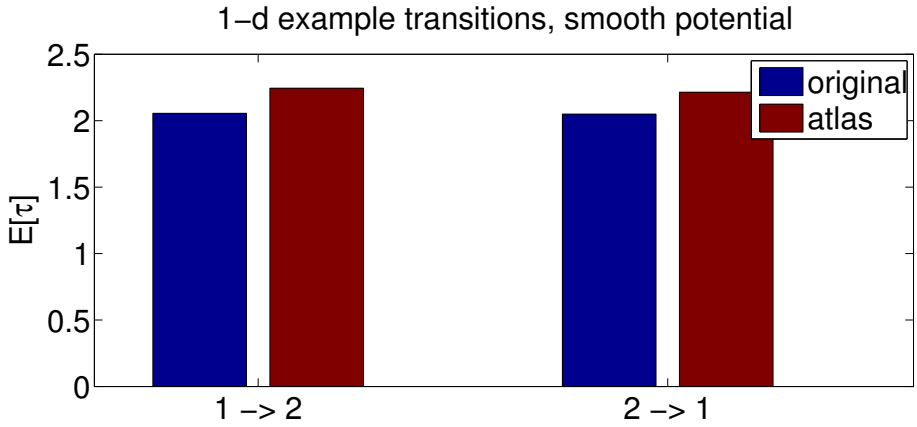

Figure 12: Comparing transition times in example 5.2.1. Region 1 is $\{x : |x| < 1/4\}$, and region 2 is $\{x : |x - 1| < 1/4\}$.

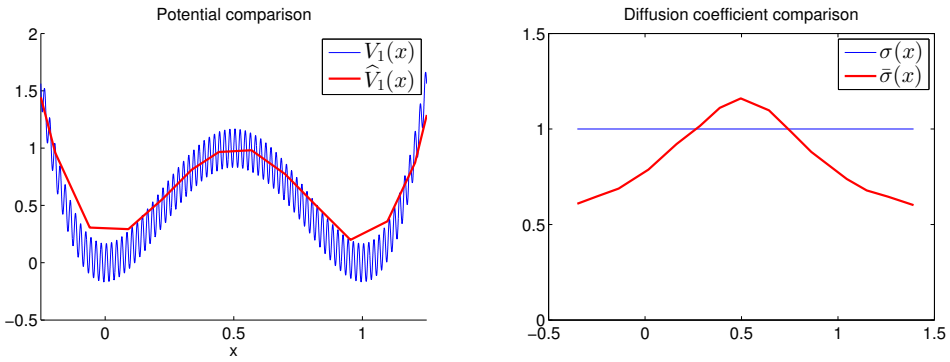

Figure 13: Left: original potential U (shown in blue) and effective potential of the ATLAS $\widehat{U}$ (shown in red). Right: comparing original diffusion coefficient (blue) with that of the ATLAS(red) with $\delta = 0.1$ in example 5.2.2.

5.2.2 Rough Potential

In order to make a more interesting example, we add high frequency ridges to the potential well to emulate microscale interactions. This example is a case where it is of interest to approximate a homogenized system which behaves like the original system above a certain temporal/spatial scale. Define

$$V_1(x) = U_1(x) + \frac{1}{6}\cos(100\pi x) \tag{86}$$

where $U_1(x)$ is defined in example 5.2.1. For our initial point set, we could again use evenly spaced grid points as in example 5.2.1. Since one might wonder if this is a "fair" input we run each grid point through the simulator for a small time $t = 0.01$ to obtain our initial point set. As long as these points have no holes of size order δ, the ATLAS will return a robust result with high probability.

Even though the new potential well is infinitely differentiable, the Lipschitz constant of the drift in this example is 625. In order to accurately simulate Brownian motion in this potential well, we decrease the time step to 0.00005. These microscale interactions are determining our timestep, and thus becoming a bottleneck for running long time simulations.

If we were to apply theorem 4.1 directly to this example, it will guarantee a relatively useless error bound on the stationary distribution (since the error bound depends on the Lipschitz constant). Instead, the way we think of theorem 4.1 applying to this problem is that there is a time scale t_0 at which the system with potential well V_1 behaves like a homogenized version with smooth potential and small Lipschitz constant. Multiscale systems of this form have been studied (see [40] and references therein), and it is known that such systems behave like an SDE with smooth parameters at a large scale. If we only observe samples at time t_0, then we can pretend our samples come from the homogenized system rather than the microscale simulator.

In this example, we learn the ATLAS using the parameters $\delta = 0.1$, $t_0 = 2\delta^2 = 0.02$, $p = 10,000$, and $\Delta t = t_0/5$. Again we can map the drift and diffusion back to the original space and compare with the true simulator. Figure 13 shows that the resulting drift is a homogenized version of the original system. The time scale the local simulator uses is 100 times larger than that of the original system. This will result in long simulations being about 100 times faster than using the original simulator.

Next we have run 10,000 long paths from the ATLAS with $\delta = 0.1$ shown above in figure 13. Figure 14 shows that again the distribution of paths is similar over multiple timescales, indicating that transition rates are preserved between states.

In this example we only show results for $\delta = 0.1$ because that is the spatial scale where it makes sense to homogenize. For smaller values of δ, the ATLAS becomes less stable as the estimated drift becomes less

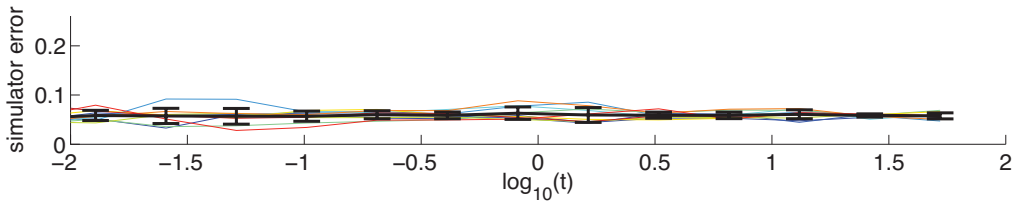

Figure 14: Comparing true simulator with the ATLAS with $\delta = 0.1$ on example 5.2.2.

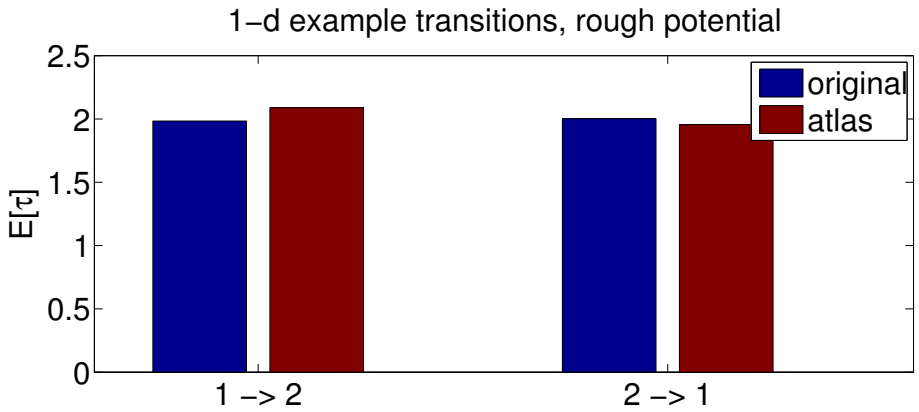

Figure 15: Comparing transition times in example 5.2.2

smooth. For larger values of δ, the macroscale features of the system begin to wash out, and the two wells merge into one.

5.3 Two Dimensional Example

5.3.1 Smooth Potential

In this example, we consider the SDE X_t in a 2-d potential well $U_2(x)$ shown below.

$$dX_t = -\nabla U_2(X_t)dt + dB_t$$

where

$$U_2(x) = -\ln\left(\exp\left(\frac{-||x-p_1||^2}{c_1}\right) + \exp\left(\frac{-||x-p_2||^2}{c_2}\right) + \exp\left(\frac{-||x-p_3||^2}{c_3}\right)\right)$$

$$p_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad p_2 = \begin{bmatrix} 1.5 \\ 0 \end{bmatrix}, \quad p_3 = \begin{bmatrix} 0.8 \\ 1.05 \end{bmatrix}, \quad c = \left[\frac{1}{5}, \frac{1}{5}, \frac{1}{6}\right]$$

The potential U_2 is chosen such that the stationary distribution is a mixture of Gaussians given by $\exp(-U_2/2)$. There are three clearly defined minima of U_2 close to p_1, p_2, p_3. The parameters of the problem were chosen in such a way that the transition regions between wells lie on different level sets of the potential (see figure 16).

Figure 17 shows a sample trajectory of the process X_t using a simple Euler-Maruyama scheme with timestep 0.005. This is the simulator given to the ATLAS algorithm. The initial point set we use is a grid spaced by 0.01, discarding points with $U_2(x) \geq 10$. Figure 16 shows an example net for $\delta = 0.2$.

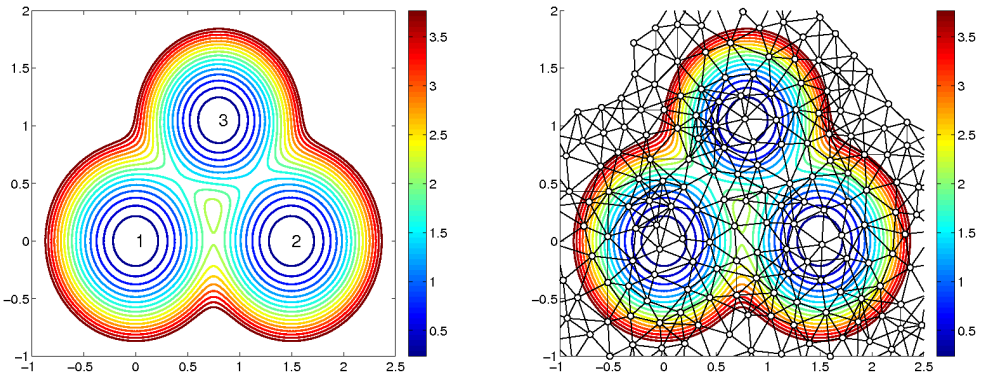

Figure 16: Left: Potential for three well example. Right: $\delta = 0.2$ net overlaid. Circles represent net points, black lines represent connections between net points.

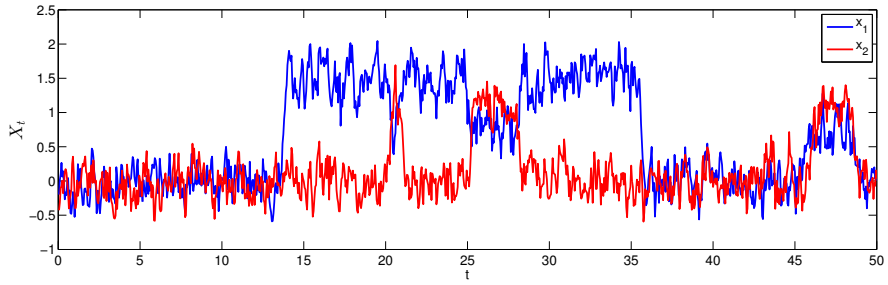

Figure 17: Sample trajectory for three well example

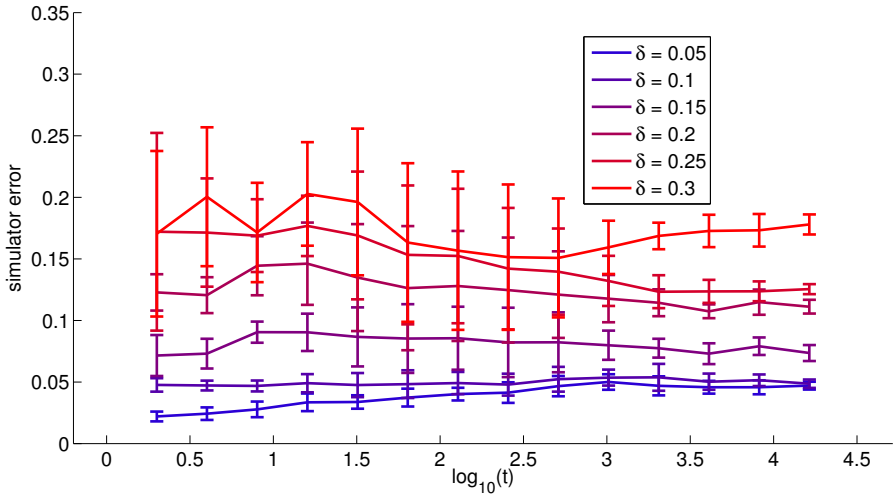

Figure 18: Comparison of ATLAS's with original simulator in the smooth three well potential from example 5.3.1.

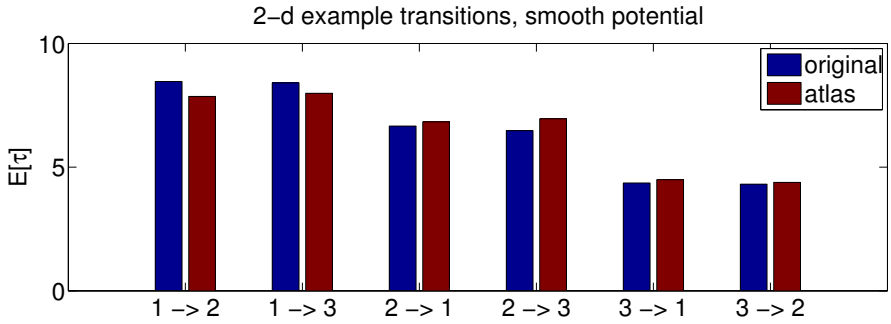

Figure 19: Comparing transition times in example 5.3.1

When generating the ATLAS in this example, we use $p = 10,000$, $t_0 = \delta^2$, $\Delta t = t_0/5$. Again the number of landmarks does not matter since LMDS will return the exact result (up to machine precision) every time. Next for each of 10 randomly chosen starting locations we run 10,000 paths from each simulator. Then we compare them using a common coarse grained net with $\delta_c = 0.2$ as in section 5.1. The output is shown in figure 18. Again we notice that the errors are small for all times, including the range of timescales where transitions occur.

In order to calculate transition times, we must first define the regions of interest. Region i will be a ball of radius $1/4$ around p_i. We will use these same regions for future examples stemming from this potential well. For a comparison of the transition times, see figure 19.

5.3.2 Rough Potential

In the next example we take $U_2(x)$ and add a fast oscillating component to simulate small scale interactions as in example 5.2.2. The new potential well is

$$V_2(x) = U_2(x) + \frac{1}{6}\cos(100\pi x_1) + \frac{1}{6}\cos(100\pi x_2). \tag{87}$$

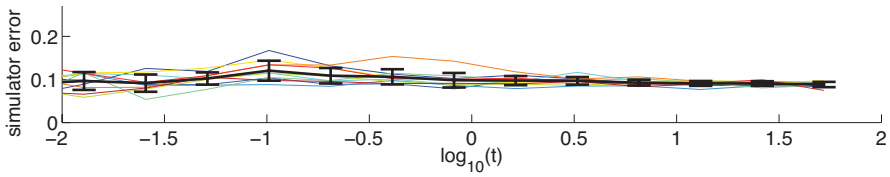

Figure 20: Comparison of original simulator with the ATLAS ($\delta = 0.2$) in the rough three well potential from example 5.3.2.

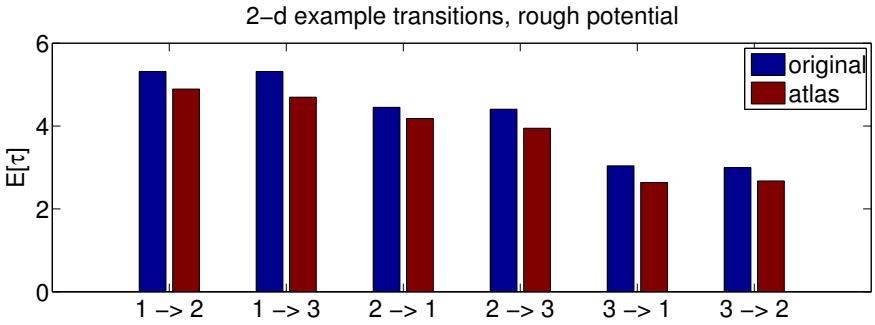

Figure 21: Comparing transition times in example 5.3.2

And again see a simulator which approximates the process X_t.

$$dX_t = -\nabla V_2(X_t)dt + dB_t \tag{88}$$

As a result of the high frequency oscillations, the the timesteps will be of size 0.00005. This example will show that our algorithm is robust to fast oscillations of the potential even in a more complicated system. In this example we will again avoid using evenly spaced points as input, and run the grid points through the simulator for a short time $t = 0.01$. These are samples we could obtain from running the original simulator for a long time, or using some kind of fast exploration technique. Again, the distribution of this point set is irrelevant as long as there are no holes of size δ.

For this system we will use $\delta = 0.2$ which will return δ nets with ≈ 230 net points. We will again use use $p = 10,000$, $t_0 = \delta^2$, $\Delta t = t_0/5$ for consistency, even though p could be chosen smaller (since δ is larger). Again, the timestep of the ATLAS is $\Delta t = 0.004$ which is over 100 times larger than the timesteps of the original simulator, and thus the ATLAS runs about 100 times faster. For the simulator comparison with this example see figure 20. Define the regions the same as in example 5.3.1. To see the transition times, see figure 21.

5.4 Random Walk on Images

Next we will embed the two dimensional three well examples from sections 5.3.1 and 5.3.2 into $D = 12,500$ dimensions. The high dimensional embedding is given by the following algorithm given a two dimensional point x:

1. Generate a mesh $\{z_j\}$ on $[-1.5, 3.5] \times [-1.5, 2.5]$ with evenly spaced grid points and spacing 0.04.
2. The output vector v at position j is 1 if $|z_j - x| < 1/2$ and 0 otherwise.

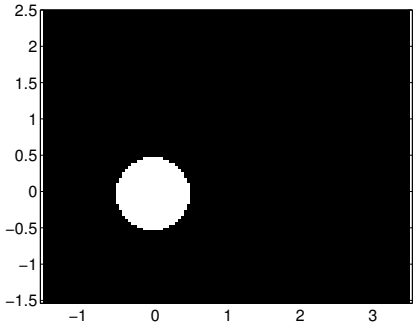

Figure 22: Circle image corresponding to the point [0,0].

See figure 22 for an example image generated by this algorithm run on the point $(0, 0)$. The natural distance to use in this space is the hamming distance, which counts the number of different entries. It induces a norm, which we call $||v||_1$ since this is the same as the 1-norm of the vector on $\mathbb{R}^D$. Given a binary vector v, we can write the "inverse" $\widetilde{x}$

$$\widetilde{x} = ||v||_1^{-1} \sum_j v_j z_j \tag{89}$$

This just averages the positions of the pixels $\{z_j\}$, which should roughly return the center of the circle in the image. Any two dimensional simulator now can be mapped to a simulator on $\mathbb{R}^D$ in the following way:

1. Given input $v \in \mathbb{R}^D$ and a time t_0, calculate the two dimensional point $\widetilde{x}$ from the approximate inverse mapping.
2. Run the 2-d simulator for time t_0 with initial condition $x_0 = \widetilde{x}$.
3. Take the output of the simulator, X_{t_0} and map it to $\mathbb{R}^D$ with the high dimensional embedding.

Next, we rescale the distance function by the constant $(0.04)^2/2$ so that the new norm is locally equivalent to the original distance. In so doing, we can continue using values of δ that made sense to us in the original space. This high dimensional mapping is nontrivial, and all the possible vectors v we could see span the entire $12,500$ dimensional space. The space can be locally approximated by a 2-d plane for a ball of radius $r < 1/2$, and so we expect the ATLAS to find the appropriate local spaces to estimate the dynamics.

5.4.1 Smooth Potential

First we will apply the high dimensional mapping to the simulator with smooth potential well U_2 from example 5.3.1. Next we start with a set of points in $\mathbb{R}^D$ which cover the known state space (the same covering set from before only mapped to $\mathbb{R}^D$). The ATLAS algorithm is given the rescaled hamming distance function for computing distances between vectors, and it is given the simulator which takes points in $\mathbb{R}^D$ and a time t_0 and returns points in $\mathbb{R}^D$. Because distances are now $12,500$ times more expensive to compute, for this example we set $p = 1000$ and $m = 20$ landmarks per point. Again keep $t_0 = \delta^2$ and $\Delta t_0 = t/5$.

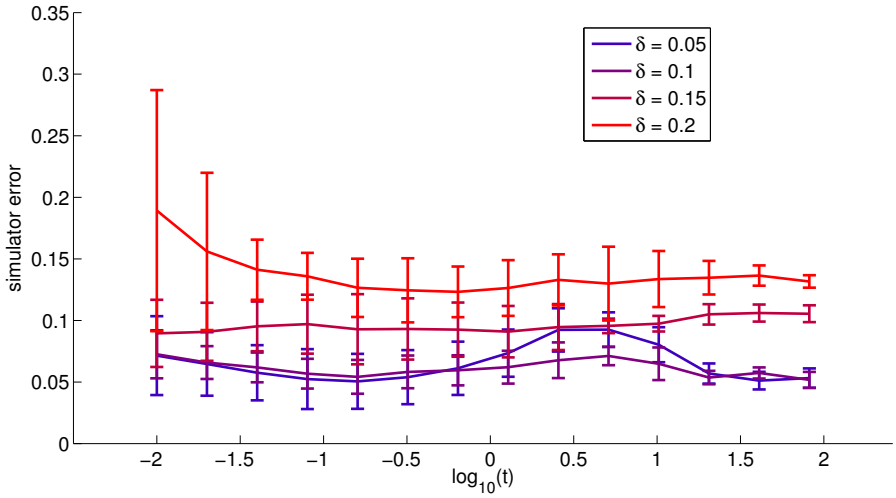

Figure 23: Comparison of ATLAS's with original simulator for example 5.4.1.

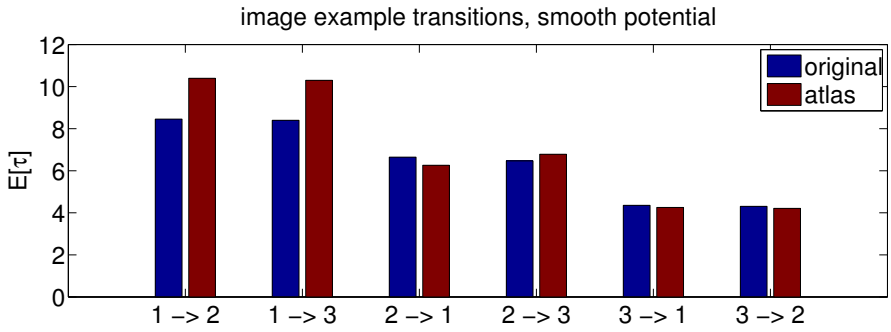

Figure 24: Comparing transition times in example 5.4.1

After constructing multiple ATLAS's for varying values of δ, we find that the distributions are well approximating the original given simulator. See figure 23 for details. The small number of samples, along with the width of the pixels limits the accuracy for small values of δ. In fact we can see that $\delta = 0.05$ returns a simulator which is worse than $\delta = 0.1$. Define the regions the same as in example 5.3.1. To see the transition times, see figure 24.

5.4.2 Rough Potential

In the next example of this paper, we will apply the high dimensional transformation to the rough potential well V_2 from example 5.3.2. Again, we give the algorithm the same set of initial points from example 5.3.2 mapped to $\mathbb{R}^D$ along with the simulator using V_2 embedded in high dimensions. In this example we use $\delta = 0.2$, $p = 2000$, $m = 40$, $t_0 = \delta^2$ and $\Delta t = t_0/5$. Again the simulation timescale of the local simulator is 100 times larger than that of the original simulator. The ATLAS has a running time which depends only on the local dimensionality of the system, and so the ambient dimension only enters in the construction phase.

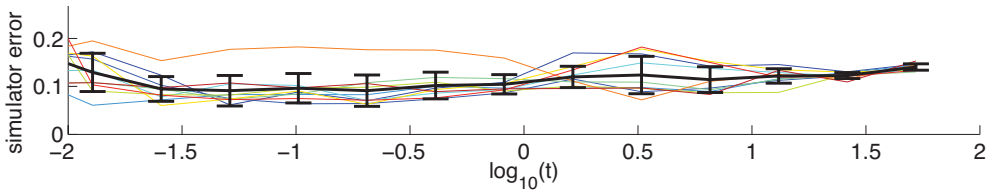

Figure 25: Comparison of the ATLAS $\delta = 0.2$ with original simulator for example 5.4.2.

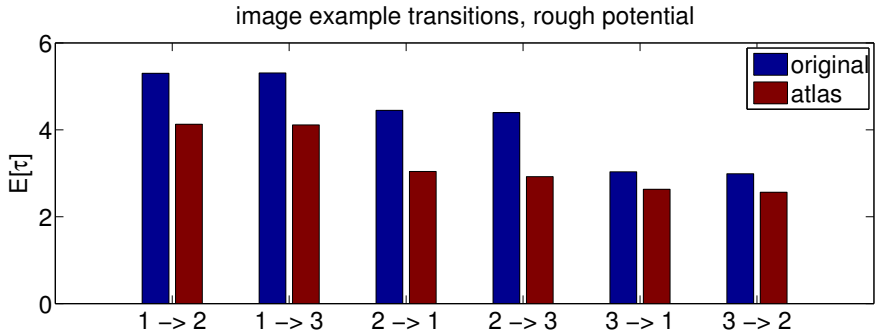

Figure 26: Comparing transition times in example 5.4.2

After simulating $10,000$ paths for each of 10 different initial conditions, we can test the simulator error (see figure 25). Because running the original simulator is very expensive for this system, we used the same original simulator samples (mapped to $\mathbb{R}^D$) for comparison as in figure 20. Define the regions the same as in example 5.3.1. To see the transition times, see figure 26.

5.5 Randomly forced string

In this example, we are given a dynamical system in the form of a random walk on functions on $[0, 1]$ with endpoints fixed at zero. These functions are represented by values on a grid of 100 evenly spaced points (including the ends). Typical functions seen as output from the simulator are shown in figure 28. The distance we will use is euclidean distance in $\mathbb{R}^{100}$, rescaled by $1/100$ to approximate the L^2 distance on functions. A single step of the simulator is done by adding a Brownian path fixed at the endpoints, then smoothing the result and renormalizing. The pseudocode is shown in figure 27.

This behavior of this system in characterized by large dwelling times near the smoothest functions (f_1 and f_2 from figure 28) with rare transitions ($10^3 - 10^4$ steps) across functions like that of f_3 in figure 28. The three constraints $f(0) = 0, f(1) = 0, ||f|| = ||f_0||$ force the functions to live on S^{97}, a 97 dimensional sphere with radius $||f_0||$. Although we expect these functions to lie near a low dimensional submanifold $\mathcal{M} \subset S^{97}$ because of the smoothing step, a single step of the simulator could take us anywhere on S^{97}; this means the outputs of our simulator are never exactly on $\mathcal{M}$. This is an important aspect of this example, as real world data typically will have small noise in the ambient space.

One can think of this simulator as a discretization of the SDE on S^{97}

$$dX_t = F(X_t)dt + \sigma(X_t)dW_t \tag{90}$$

Function Simulator

$f = \mathcal{S}(f)$

$\%$ *simulate Brownian bridge*
$W = \text{cumsum}(\text{randn}(1{,}100))$
$W = W - W(1)$
$W = W - x * W(100)$

$\%$ *Add bridge to f, smooth and renormalize*
$f = f + (1/100) * W$
$f = \text{smooth}(f)$
$f = f * (f_{\text{norm}}/\text{norm}(f))$

Figure 27: Pseudocode for a single step of the simulator used in example 5.5. f_{norm} is a fixed constant equal to the norm of $\sin(\pi x)$. The function smooth is MATLAB's default smoothing algorithm.

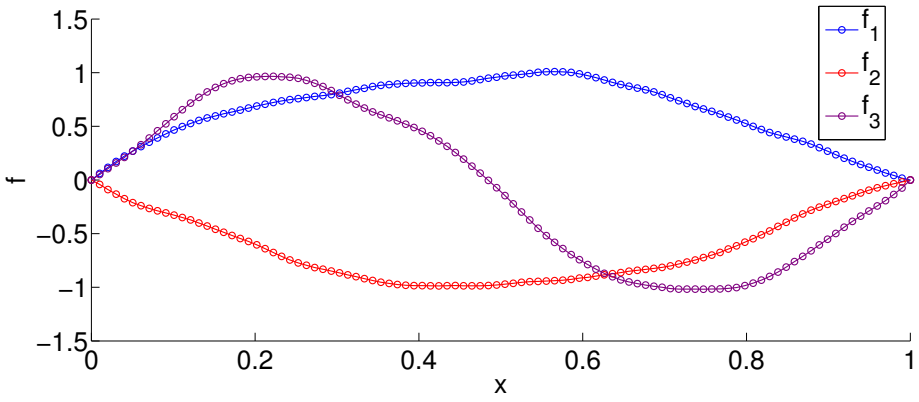

Figure 28: Three typical outputs of the simulator from section 5.5.

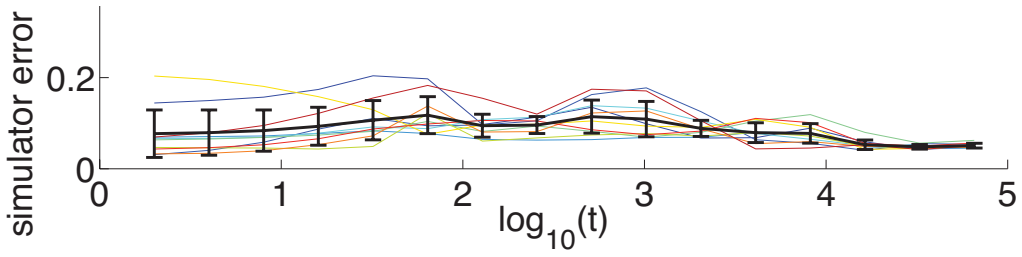

Figure 29: Comparison of the ATLAS with original simulator for example 5.5.

For an appropriate choice of F, σ. One can also think of this as a discretization scheme for a stochastic partial differential equation (SPDE) of the form

$$\frac{\partial}{\partial t} f_t = \frac{\partial^2}{\partial x^2} f_t + b(f_t) + \sum_{j=1}^{\infty} g_j(f_t) dW_t^j \tag{91}$$

for an appropriate choice of drift b and orthogonal functions $\{g_j\}$. One can think of (91) as an infinite dimensional analogue to (90) with each coordinate $X_t^j = \langle f_t, g_j \rangle$ being driven by a one dimensional brownian motion.

In order to generate the ATLAS, first we must generate an initial sampling of the space. In order to do this, we start with $50,000$ renormalized Gaussian vectors, the uniform distribution on S^{97}. Next we want to "heal" these samples by running them through the simulator. One can see with some observation that 250 steps is large enough that the noise is killed; samples with 250 steps of "healing" are similar to those with 500 steps of "healing".

Next we wish to select parameters δ, t_0. We expect that the system may be homogenized at a time scale of $t_0 = 250$ steps for the following reasons: t_0 is an order of magnitude below the scale of major events of the system, t_0 is an order of magnitude above the scale of the noise (since even the most noisy inputs have been smoothed by time t_0). The parameter δ is closely tied to the choice of t_0. We measure the average distance moved by paths of length t_0 starting from our healed samples to be $0.3 = \delta$. Next we choose the minor parameters p, m, d. In these experiments, we use $p = 5000$ and $m = 40$. As discussed in section 3.2, we can choose d based upon the singular values obtained through LMDS. Choosing a cutoff of $(\delta/4)^2$ for the eigenvalues yields $d = 3$ over 99% of the time. Using $d = 3$ and comparing with the original simulator in the usual way yields figure 29.

In general, it is better to overestimate d than underestimate; underestimating d may lose important degrees of freedom causing failure, while overestimating d will only affect the computational cost mildly. In fact the algorithm is robust to the choice of d, provided d is large enough to capture the important degrees of freedom. See figure 31 to see results for varying values of the choice of d.

Next we wish to compare the transition times between states. In order to do this, we define region 1 to be a ball of radius 1/4 around $\sin(\pi x)$, and region 2 to be a ball of radius 1/4 around -$\sin(\pi x)$. To see a comparison of transition times between these regions, see figure 30.

The ATLAS constructed for this example again captures the important aspects of the original simulator. The ATLAS is again faster in this example due to two factors: decreased dimensionality and increased timestep. The dimensionality of the ATLAS is 3 as compared to the original 100, and the timestep of the simulator is equivalent to 50 of the original steps (250/5 since $\Delta t = t_0/5$).

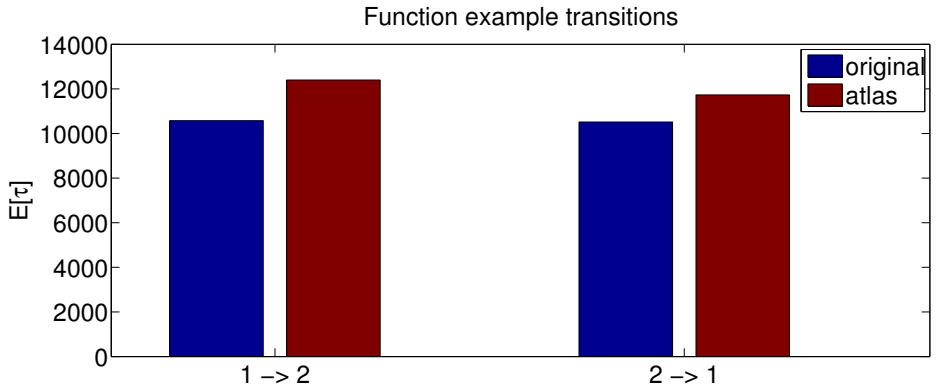

Figure 30: Comparing transition times in example 5.5

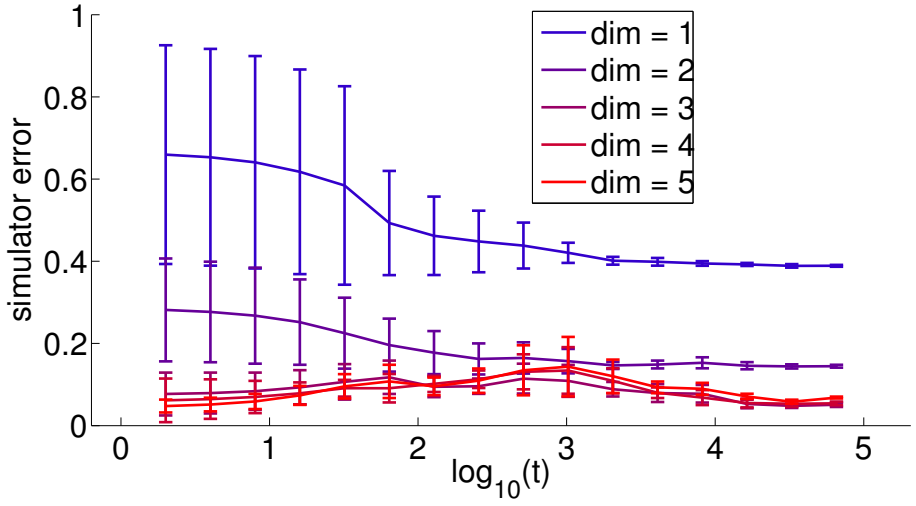

Figure 31: Comparison for example 5.5 varying d, the dimension of the ATLAS.

5.6 Chaotic ODE system

In many real world systems, noise arises from ensembles of deterministic chaotic processes. In this example we apply our algorithm to a multiscale ODE driven by small scale deterministic chaos. Consider the set of multiscale ODEs with a scale parameter ε:

$$\begin{cases} \dot{X}_t^\varepsilon = \varepsilon f(X_t^\varepsilon) + g(Y_t), & X_0^\varepsilon = x \\ \dot{Y}_t = h(Y_t) & Y_0 = y \end{cases} \tag{92}$$

Systems of this form (although slightly more general) are studied in [53]. Suppose the dynamics for Y_t alone have an invariant measure μ, and $\mathbb{E}_\mu[f] = O(\varepsilon)$. Then the system (92) behaves like the SDE

$$dX_s = b(X_s)ds + \sigma(X_s)dB_s \tag{93}$$

on the timescale $s = \varepsilon t$ in the limit as $\varepsilon \to 0$. For fixed ε, such systems are difficult to simulate directly due to the timescale separation.

We start by choosing functions f, g, h and scale parameter ε. Start by choosing Y_t to be the Lorenz '96 system with 80 dimensions and $F = 8$ (thus fixing h). Each coordinate $Y_i(t)$ is governed by equation 94, where indices wrap around (so $Y_{-1} = Y_{79}, Y_0 = Y_{80}, Y_{81} = Y_1$).

$$\dot{Y}_i = -Y_{i-2}Y_{i-1} + Y_{i-1}Y_{i+1} - Y_i + F \tag{94}$$

We fix $\varepsilon = 0.01$ and let $f(X_t)$ be the cartesian coordinate version of the system in equations 95,96:

$$\dot{r} = -(r - 3/4)(r - 3/2)(r - 2) \tag{95}$$

$$\dot{\theta} = r - 3/2 \tag{96}$$

Last, we choose $g = [g_1(y), g_2(y)]$ to be

$$g_1(y) = \frac{1}{320}\sum_{i \in I_1} y_i - 0.2925 \quad , \quad g_2(y) = \frac{1}{320}\sum_{i \in I_2} y_i - 0.2925 \tag{97}$$

with $I_1 = [1 : 10, 21 : 30, 41 : 50, 61 : 70]$ and $I_2 = I_1^c$. This choice of g was made in order that g is approximately mean zero and variance 1 with respect to the measure μ. The system is then solved using the Runge-Kutta method with timestep 0.05. Last, we multiply Y_t by a small constant 10^{-4} in a post processing stage in order that these directions do not overpower the interesting states of the system (the first two coordinates). We use 10^{-4} since a typical value of each coordinate of $Y_t \approx 10$, there are 80 variables, so this ensures the norm of Y_t is lower than ε. The system then behaves like

$$d\bar{X}_s = f(\bar{X}_s) + dB_s \tag{98}$$

on the timescale $s = \varepsilon t$. The ODE system associated with f (i.e. without the stochastic term dB_s) has two stable attracting limit cycles, one at $r = 3/4$ and the other at $r = 2$. A slope field for f is shown, along with a sample path of X_t, in figure 32.

We cannot apply directly our algorithm to the above, since multiple runs will yield the same result (and have covariance zero). For this reason we will add small noise (10^{-5} times a random normal) to the initial condition we give as input, and pretend our simulator is of the form (98). Because of the chaos in the system, this small perturbation propagates quickly through the system, yielding us a different "realization" of the chaos Y_t.

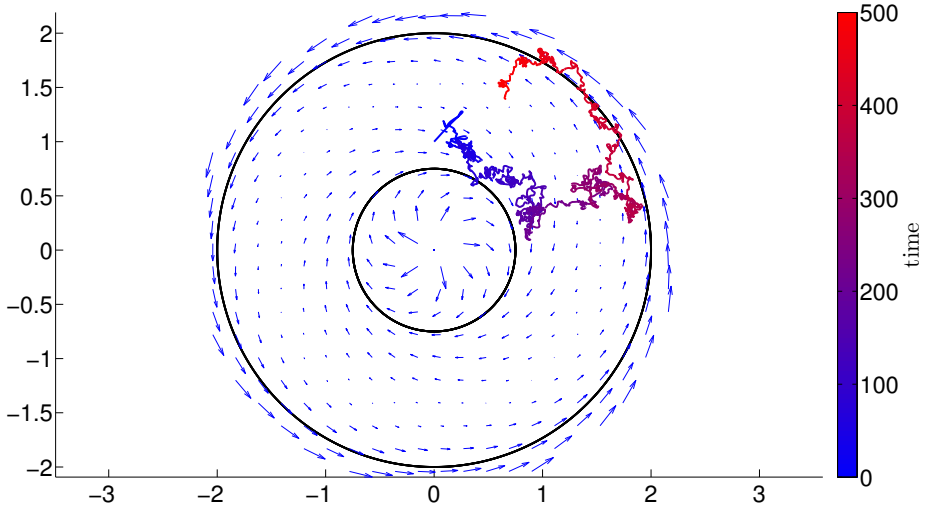

Figure 32: Colored path: example simulation for example 5.6. Black circles represent the stable limit cycles, and blue vectors represent f at that location.

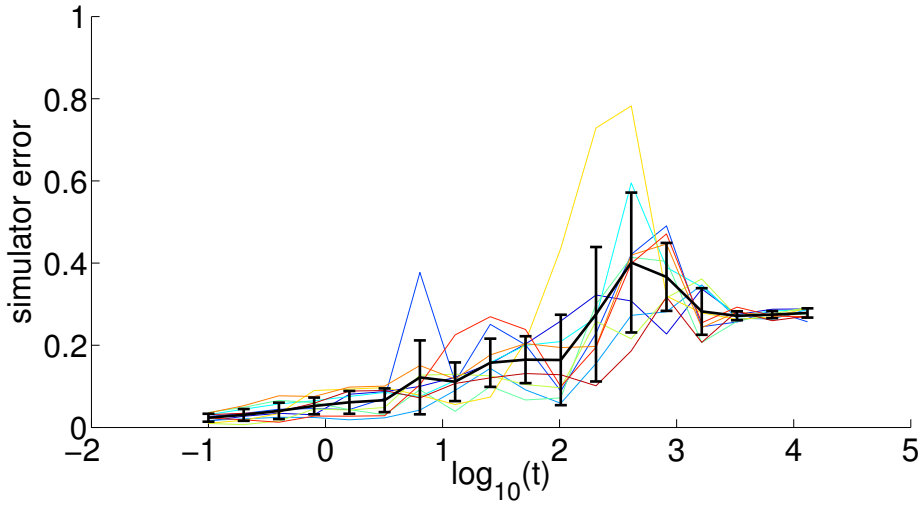

Figure 33: Comparison of the ATLAS with original simulator for example 5.6.

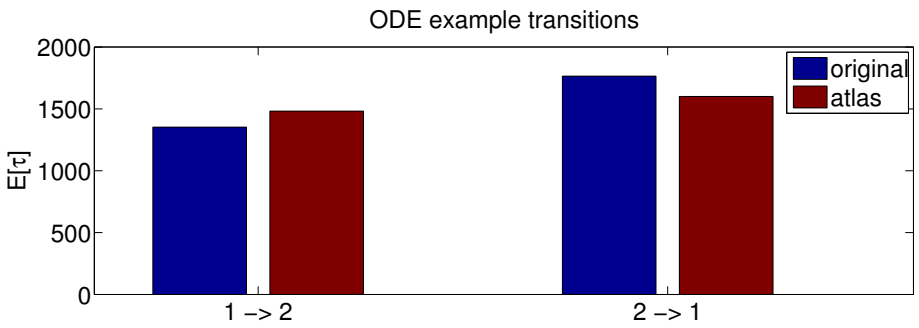

Figure 34: Comparing transition times in example 5.6

Next we need to generate a large set of input points for the δ-net. In this case we use 82 dimensional random normal vectors with the last 80 directions multiplied by 10^{-4} (in order to be similar in size to typical outputs of the simulator). We run these random samples through the simulator for a short time in order to "heal" them so that the points given to us are actual outputs of the true simulator we are given. Next we must choose $t_0 = O(\varepsilon^{-1})$, since this is the timescale on which the SDE dynamics occur. The exact value we use is $t_0 = 25 = (1/4)\varepsilon^{-1}$. In this amount of time, the system travels roughly $\delta = 0.3$. The ATLAS we construct compared with the true simulator is shown in figure 33. Here we choose region 1 to be everywhere $r < 1$, and region 2 to be everywhere $r > 7/4$. To see a comparison of the transition times between these regions see figure 34. Notice that the rates of transition between states are accurate, but the comparison is not very accurate for intermediate times. This is because the speed at which the system travels around the limit cycles is not well approximated.

6 Extensions and Discussion

There are many open problems related to this work, some of which we mention here.

Theorem 4.1 reveals that the local learning algorithm works well on compact SDEs with Lipschitz drift and diffusion. We consider only bounded domains in the proof to make thing simpler, although the same framework can be applied to the unbounded case with tight transition density. In this case, one has to worry about parts of the space which are unexplored, but seldomly reached. Indeed we see that some of our examples have unbounded state spaces, and the algorithm performs as desired.

The framework we introduced may be generalized to richer families of local simulators, enabling the approximation of larger classes of stochastic systems. Proving large time accuracy may be difficult for such systems, so it is an open problem how much one is allowed to change these local simulators. Many molecular dynamics (MD) systems remember the velocity of atoms and so do not follow an SDE of the form (18) which is memoryless. A subject of ongoing research is to use more complex models locally to be able to capture dynamics of typical MD systems.

Another subject of future work is efficient computation of the function G, which is the inverse MDS mapping. One can always approximate this function up to order δ via a piecewise constant function (returning the chart center). In some cases, such as when ρ is the root mean square distance (RMSD), it is possible to create an inverse mapping which has error of order δ^2 via local linear approximations, using ideas from [2].

Using the ATLAS as a basis for generating samples from the stationary distribution is useful for quickly computing diffusion maps for these systems. A subject of interest is to understand how the errors made by the ATLAS propagate through diffusion maps. How similar do diffusion maps look generated by samples from the ATLAS as compared to diffusion maps generated directly from the original simulator?

In some problems, choosing δ and t_0 is difficult. Another subject of ongoing research is a robust way of choosing these parameters based on short simulations. For simplicity in this paper we have assumed that δ

and t_0 is constant for each $k \in \Gamma$, but it is possible to have these parameters depend on the location y_k (and perhaps statistics of short sample paths).

Last but not least, this construction as described here still requires a large number of steps to sample rare events and reach stationarity, i.e. it does not address the problem of accelerating the sampling of rare events or overcoming energy barriers. In many important applications, e.g. molecular dynamics, such barriers force the simulations to be extremely long (e.g. $10^{12} - 10^{14}$ time-steps is common). The point of this work is to produce a simulator that is much faster (in real world time) than the original fine scale simulator. It is important to note that any of the many techniques developed over the years to attempt to overcome this problem may be used in conjunction with our construction, i.e. it can be run on our ATLAS, instead of the original expensive fine scale simulator. This yields a double gain in simulation speed, combining the gains of a faster simulator with those of an importance sampler that efficiently samples rare events.

7 Acknowledgement

The authors gratefully acknowledge partial support from NSF CAREER DMS-0847388, NSF CHE-1265920 and ONR N00014-12-1-0601. We thank J. Mattingly for useful discussions, and Y. Kevrekidis for introducing us to the questions around model reduction and equation free methods, and many discussions on these topics over the years. Finally, we would like to thank the anonymous reviewer, whose comments on the first version of this manuscript greatly helped us improve the presentation and readability of this work.

References

[1] I. Timofeyev A. J. Majda and E. Vanden-Eijnden. A mathematical framework for stochastic climate models. *Comm. Pure App. Math.*, 8(54):891–974, 2001.

[2] William K. Allard, Guangliang Chen, and Mauro Maggioni. Multi-scale geometric methods for data sets II: Geometric multi-resolution analysis. *Applied and Computational Harmonic Analysis*, 32(3):435–462, 2012. (submitted:5/2011).

[3] A. C. Antoulas, D. C. Sorensen, and S. Gugercin. A survey of model reduction methods for large-scale systems. *Contemporary Mathematics*, 280:193–219, 2001.

[4] A Berezhkovskii and Attila Szabo. One-dimensional reaction coordinates for diffusive activated rate processes in many dimensions. *J Chem Phys*, 122(1):014503, 2005.

[5] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104. ACM, 2006.

[6] P. Billingsley. *Convergence of Probability Measures*. Wiley-Interscience Publication, 1999.

[7] Groenen P. Borg, I. *Modern Multidimensional Scaling: theory and applications*. 2005.

[8] G. Bowman, K. Beauchamp, G. Boxer, and V. Pande. Progress and challenges in the automated construction of Markov state models for full protein systems. *The Journal of chemical physics*, 131(12):124101, 2009.

[9] M. Brand. Charting a manifold. In *Advances in neural information processing systems*, pages 961–968, 2002.

[10] SS Cho, Y Levy, and PG Wolynes. P versus Q: Structural reaction coordinates capture protein folding on smooth landscapes. *P Natl Acad Sci Usa*, 103(3):586–591, 2006.

[11] R. Coifman, I. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler. Diffusion maps, reduction coordinates, and low dimensional representation of stochastic systems. *Multiscale Model Sim*, 7(2):842–864, 2008.

Notation and Definitions

Notation:

- For f, g functions on the same domain, we say that $f = O(g)$ if there is a constant $C > 0$ such that for all x in the domain of f and g we have $f(x) \leq Cg(x)$.
- If $F : M \to N$ is a measurable function and μ is a measure on M then we define the push forward measure of μ through F for any $A \subset N$ by $F_*\mu(A) = \int_M \mathbb{1}_A(F(x))d\mu(x)$.

Definitions:

- $\Gamma = \{y_k\}$: Set of points in the net.
- $i \sim j$: Neighbor connections on net indices (3.1).
- Φ_i: Mapping created by LMDS associated with chart i, see section 3.2.
- $\mathcal{A}$: Atlas constituted by the collection of $\{(x, i)|x \in \mathbb{R}^d, |x| < 2\delta, i \in \Gamma\}$.
- $S_{i,j}$: Chart transition map $S_{i,j}(x) = (x - \mu_{i,j})T_{i,j} + \mu_{j,i}$, see equation (8).
- i': Shorthand denoting the new chart index chosen starting from $(x, i) \in \mathcal{A}$, see equation (14).
- W: Wall function which keeps the ATLAS within 2δ of the chart center, see equation (11).
- q: density (with respect to volume measure on $\mathcal{M}$) of the stationary distribution for the original process Y_t.
- μ: stationary distribution of the ATLAS process Z_k.
- $\widehat{q}$: 'almost stationary' distribution for $\widehat{X}_t$ (see (16)).
- $C^k, ||\cdot||_{C^k}$: Hölder space of order k.

Figure 35: List of notation and definitions used for reference

[12] R.R. Coifman and S. Lafon. Diffusion maps. *submitted to Applied and Computational Harmonic Analysis*, 2004.

[13] P. Das, M. Moll, H. Stamati, L.E. Kavraki, and C. Clementi. Low-dimensional, free-energy landscapes of protein-folding reactions by nonlinear dimensionality reduction. *P Natl Acad Sci Usa*, 103(26):9885–9890, 2006.

[14] P. Das, M. Moll, H. Stamati, L.E. Kavraki, and C. Clementi. Low-dimensional, free-energy landscapes of protein-folding reactions by nonlinear dimensionality reduction. *P.N.A.S.*, 103(26):9885–9890, 2006.

[15] P Deuflhard, W Huisinga, A Fischer, and Christof Schutte. Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Linear Algebra Appl*, 315(1-3):39–59, 2000.

[16] W. E, W. Ren, and E. Vanden-Eijnden. Simplified and improved string method for computing the minimum energy paths in barrier-crossing events. *J Chem Phys*, 126(16):164103, 2007.

[17] Emery. *Stochastic Calculus in Manifolds*. 1989.

[18] A. Faradjian and R. Elber. Computing time scales from reaction coordinates by milestoning. *The Journal of Chemical Physics*, 120(23), 2004.

[19] M. Maggioni G. Chen, A. Little. Multi-resolution geometric analysis for data in high dimensions. In Travis D. Andrews, Radu Balan, John J. Benedetto, Wojciech Czaja, and Kasso A. Okoudjou, editors, *Excursions in Harmonic Analysis, Volume 1*, Applied and Numerical Harmonic Analysis, pages 259–285. Birkhäuser Boston, 2013.

[20] A. Stuart G. Pavliotis. *Multiscale Methods*. Springer, 2008.

[21] C. Gear, T. Kaper, I. Kevrekidis, and A. Zagaris. Projecting to a slow manifold: Singularly perturbed systems and legacy codes. *SIAM Journal on Applied Dynamical Systems*, 4(3):711–732, 2005.

[22] AC Gilbert. A comparison of multiresolution and classical one-dimensional homogenization schemes. *Applied and Computational Harmonic Analysis*, 5(1):1–35, 1998.

[23] M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73:123–214, 2011.

[24] P. Hanggi, P. Talkner, and M. Borkovec. Reaction-rate theory - 50 years after kramers. *Rev Mod Phys*, 62(2):251–341, 1990.

[25] U. Hornung. *Homogenization and Porous Media*. Springer, 1997.

[26] E. Hsu. *Stochastic Analysis on Manifolds*. American Mathematical Society, 2002.

[27] W Huisinga, S Meyn, and Christof Schutte. Phase transitions and metastability in Markovian and molecular systems. *Ann Appl Probab*, 14(1):419–458, 2004.

[28] W Huisinga, Ch. Schütte, and A Stuart. Extracting macroscopic stochastic dynamics. *Comm. Pure Appl. Math.*, 56(2):234 – 269, 2003.

[29] I. Kevrekidis, C. Gear, J. Hyman, and P. Kevrekidid. Equation-free, coarse-grained multiscale computation: Enabling mocroscopic simulators to perform system-level analysis. *Communications in Mathematical Sciences*, 1(4):715—762, 2003.

[30] N. Krylov. *Lectures on Elliptic and Parabolic Equations in Hölder Spaces*. American Mathematical Society, 1996.

[31] A. Little, M. Maggioni, and L. Rosasco. Multiscale geometric methods for data sets I: Multiscale SVD, noise and curvature. Technical report, MIT-CSAIL-TR-2012-029/CBCL-310, MIT, Cambridge, MA, September 2012.

[32] A.V. Little, Y.-M. Jung, and M. Maggioni. Multiscale estimation of intrinsic dimensionality of data sets. In *Proc. A.A.A.I.*, 2009.

[33] M. Maggioni, S. Minsker, and N. Strawn. Geometric Multi-Resolution dictionary and manifold learning: Non-asymptotic bounds, noise, and rough sets. *in preparation*, 2013.

[34] J. Mattingly, A. Stuart, and M. Tretyakov. Convergence of numerical time-averaging and stationary measures via Poisson equations. *SIAM Journal on Numerical Analysis*, 48(2):552–577, 2010.

[35] B. C. Moore. Principal component analysis in linear system: Controllability, observability and model reduction. *IEEE Transactions on Automatic Control*, (26):17–32, 1981.

[36] B. Nadler, S. Lafon, R.R. Coifman, and I.G. Kevrekidis. Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *Appl. Comp. Harm. Anal.*, 21:113–127, 2006.

[37] F. Noé, C. Schutte, E. Vanden-Eijnden, L. Reich, and T.R. Weikl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *P Natl Acad Sci Usa*, 106(45):19011–19016, 2009.

[38] B. Øksendal. *Stochastic Differential Equations*. Springer, 2003.

[39] V. Pande, K. Beauchamp, and G. Bowman. Everything you wanted to know about Markov state models but were afraid to ask. *Methods*, 52(1):99–105, 2010.

[40] GA Pavliotis and AM Stuart. Parameter estimation for multiscale diffusions. *Journal of Statistical Physics*, 127(4):741–781, 2007.

[41] R. Penrose. On best approximate solutions of linear matrix equations. In *Proceedings of the Cambridge Philosophical Society*, volume 52, pages 17–19. Cambridge Univ Press, 1956.

[42] B. Peters and B. Trout. Obtaining reaction coordinates by likelihood maximization. *J Chem Phys*, 125(5):054108, 2006.

[43] B. Qi, S. Muff, A. Caflisch, and A. Dinner. Extracting physically intuitive reaction coordinates from transition networks of a β-sheet miniprotein. *J Phys Chem*, 114:6979–6989, 2010.

[44] W Ren, E Vanden-Eijnden, P Maragakis, and WN E. Transition pathways in complex systems: Application of the finite-temperature string method to the alanine dipeptide. *J Chem Phys*, 123(13):134109, 2005.

[45] M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi. Determination of reaction coordinates via locally scaled diffusion map. *J. Chem. Phys.*, (134):124116, 2011.

[46] L. Saul and S. Roweis. Think globally, fit locally: unsupervised learning of low dimensional manifolds. *J. Mach. Learn. Res.*, 4:119–155, December 2003.

[47] T. Shardlow and A. M. Stuart. A perturbation theory for ergodic properties of Markov chains. *SIAM J. Num. Anal*, 37, 2000.

[48] V. De Silva and J. Tenenbaum. Sparse multidimensional scaling using landmark points. Technical report, Technical report, Stanford University, 2004.

[49] S.Larson, C. Snow, M. Shirts, et al. Folding@ home and genome@ home: Using distributed computing to tackle previously intractable problems in computational biology. 2002.

[50] ND Socci, JN Onuchic, and PG Wolynes. Diffusive dynamics of the reaction coordinate for protein folding funnels. *J Chem Phys*, 104(15):5860–5868, 1996.

[51] D. Stroock. *Partial Differential Equations for Probabilists*. Cambridge University Press, 2008.

[52] A. Szabo, K. Schulten, and Z. Schulten. 1st passage time approach to diffusion controlled reactions. *J Chem Phys*, 72(8):4350–4357, 1980.

[53] E. Vanden-Eijnden. Fast communications: Numerical techniques for multi-scale dynamical systems with stochastic effects. *Communications in Mathematical Sciences*, 1(2):385–391, 2003.

[54] E. Vanden-Eijnden and M. Venturoli. Revisiting the finite temperature string method for the calculation of reaction tubes and free energies. *J Chem Phys*, 130(19):194103, 2009.

[55] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *Compressed Sensing, Theory and Applications*, Chapter 5, 2012.

[56] Romand Vershynin. Introduction to the non-asymptotic analysis of random matrices. 2010.

[57] E Weinan, Bjorn Engquist, Xiantao Li, Weiqing Ren, and Eric Vanden-Eijnden. Heterogeneous multiscale methods: a review. *Commun Comput Phys*, 2(3):367–450, 2007.

[58] M. Pollicott & M. Yuri. *Dynamical Systems and Ergodic Theory*. London Mathematical Society, 1998.

[59] H. Zha Z. Zhang. Principal manifolds and nonlinear dimension reduction via local tangent space alignment. *CoRR*, cs.LG/0212008, 2002.

[60] G. Da Prato & J. Zabczyk. *Ergodicity for Infinite Dimensional Systems*. Cambridge University Press, 1996.

[61] W. Zheng, M. Rohrdanz, and C. Clementi. Rapid exploration of configuration space with diffusion-map-directed molecular dynamics. *The Journal of Physical Chemistry B*, 2013.

[62] W. Zheng, M. A. Rohrdanz, M. Maggioni, and C. Clementi. Polymer reversal rate calculated via locally scaled diffusion map. *J. Chem. Phys.*, (134):144108, 2011.